

Beyond Time and Iteration: Alternative Orderings of Change in Physics, Information, and Experience

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 24, 2025

Modern conversation about "what time really is" often oscillates between two intuitions: time as a continuous dimension and iteration as a discrete update count. But physics only adjudicates between ontologies when they generate different predictions, and many "time vs iteration" debates collapse into re-labeling. This paper surveys a wider possibility space: ways that "order" can arise without fundamental time, without fundamental iteration, or with both only as emergent bookkeeping. We examine correlation-first (relational) accounts where clocks are subsystems and time is a mapping between degrees of freedom; causal-order accounts where the primitive is influence rather than a global parameter; thermodynamic and record-based accounts where the arrow is built from low-entropy boundary conditions and the accumulation of stable traces; and gauge-time accounts where different parameterizations describe the same invariant content. We then widen further to phase ordering, scale ordering (renormalization flow), inference/description ordering, and agency/decision ordering—each offering a distinct way to define "before/after" without presuming a universal flowing time. Throughout, we emphasize operational discriminators: what each ontology must claim, what it can avoid claiming, and what empirical or conceptual signatures would distinguish it. The goal is a disciplined map of options, not a verdict.

Keywords: time, iteration, relational time, causal order, entropy arrow, records, block universe, gauge time, reparameterization, emergent time, discreteness, Lorentz invariance, dispersion, renormalization group, phase ordering, computational depth, inference ordering, agency, clocks, observables, ontology.

A collaboration with GPT-5.2-Thinking. 76 pages. CC4.0.

Outline

1. Introduction: Why “Time vs Iteration” Is Too Narrow
 - 1.1 The definitional trap: re-labeling versus new physics
 - 1.2 What physics can and cannot decide
 - 1.3 The paper’s aim: a taxonomy of orderings of change
2. What Counts as an Ordering in Physics
 - 2.1 Parameters, observables, and invariants
 - 2.2 Global order, local order, and partial order
 - 2.3 The role of clocks as physical subsystems
 - 2.4 Operationalism: “order” as measurable correlation structure
3. Time as a Coordinate Versus Time as Becoming
 - 3.1 The block-universe stance: time without flow
 - 3.2 Becoming/flow as an additional claim
 - 3.3 The “experience gap”: why flow feels unavoidable
 - 3.4 What would it mean for flow to be physical
4. Iteration: Discreteness, Update Rules, and Hidden Assumptions
 - 4.1 Iteration as mere reparameterization
 - 4.2 Iteration as fundamental discreteness (a true tick)
 - 4.3 Iteration as computation: update rules and resource bounds
 - 4.4 Where iteration hides a preferred frame or structure
5. Correlation-First Ordering (Relational Time)
 - 5.1 “Time from clocks”: internal time variables
 - 5.2 Multiple clocks and the problem of clock choice
 - 5.3 When relational time mimics ordinary time perfectly
 - 5.4 Conceptual payoffs and costs: realism about parameters
6. Causal Order as the Primitive
 - 6.1 Partial orders: influence without global simultaneity
 - 6.2 Reconstructing geometry from causality
 - 6.3 Locality, signaling, and causal horizons
 - 6.4 How causal order differs from discrete iteration
7. Entropy and the Thermodynamic Arrow
 - 7.1 Low-entropy boundary conditions and “direction”
 - 7.2 Why records, memories, and prediction align
 - 7.3 Irreversibility: fundamental or emergent
 - 7.4 The difference between arrow and parameter

- 8. Records Ontology: The World as Consistent Traces
 - 8.1 Records as localized correlations
 - 8.2 “Past” as an inference from present constraints
 - 8.3 The stability problem: why records persist
 - 8.4 What records can explain and what they cannot
- 9. Gauge Time and Reparameterization Symmetry
 - 9.1 Time as gauge: why parameter choice can be unphysical
 - 9.2 Fixing time by choosing a clock: what becomes real
 - 9.3 Empirical equivalence: when rival stories are the same theory
 - 9.4 How “iteration” often reduces to gauge choice
- 10. Phase Ordering: Cycles, Angles, and Return
 - 10.1 Periodic variables as order parameters
 - 10.2 Cyclic cosmologies and effective cyclic dynamics
 - 10.3 When a phase variable replaces linear time
 - 10.4 Distinguishing phase-order from record-order
- 11. Scale Ordering: Renormalization as a Kind of “Before/After”
 - 11.1 Ordering by description scale, not moments
 - 11.2 Effective theories and emergence as a flow
 - 11.3 What scale-order explains about stability and universality
 - 11.4 The limits: why scale-order is not a full substitute for time
- 12. Inference and Computational Depth as Ordering
 - 12.1 Ordering by derivability, proof depth, or compressibility
 - 12.2 Description length versus physical process
 - 12.3 When “complexity” functions like an arrow
 - 12.4 Risks: confusing epistemic order with ontic order
- 13. Agency and Decision Ordering
 - 13.1 Control loops: order defined by information availability
 - 13.2 Belief update sequences and the arrow of learning
 - 13.3 Multi-agent consistency: when orders conflict
 - 13.4 What agency-order can illuminate about experience
- 14. Empirical and Conceptual Discriminators
 - 14.1 What would count as evidence for fundamental discreteness
 - 14.2 Lorentz invariance, dispersion, anisotropy, and noise floors
 - 14.3 Unitarity, reversibility, and resource-bounded physics

- 14.4 When theories are empirically equivalent but metaphysically distinct
- 14.5 The “minimal claim” ladder: from operational order to ontology

15. Synthesis: A Map of Options

- 15.1 A decision tree: which commitments imply which structures
- 15.2 Compatibility chart: which frameworks can coexist
- 15.3 Where “time” and “iteration” sit inside the larger space
- 15.4 A pragmatic stance: disciplined pluralism and falsifiable edges

16. Conclusion: How to Speak About Order Without Overclaiming

- 16.1 Restating the central result: many orderings, few discriminators
- 16.2 Practical takeaway for theorists and communicators
- 16.3 The open frontier: new tests, new formalisms, cleaner language

17. References

- 17.1 Relational and emergent time
- 17.2 Causal set and causal order approaches
- 17.3 Thermodynamic arrow and low-entropy boundary work
- 17.4 Records, decoherence, and information-theoretic treatments
- 17.5 Gauge time, parameterization invariance, and canonical gravity
- 17.6 Discreteness tests, Lorentz invariance constraints, dispersion searches
- 17.7 Renormalization group and scale as ordering
- 17.8 Complexity, computational depth, and inference-based arrows
- 17.9 Philosophy of time, flow, and operationalism

Art

1. Introduction: Why “Time vs Iteration” Is Too Narrow

The question “Are we in time or in iteration?” sounds like it offers two clean, rival ontologies: a continuous dimension (time) versus a discrete update counter (iteration). But in practice the binary often collapses under scrutiny. Physics is not a referee of metaphors; it is a discipline for extracting invariant, testable structure from experience. Many disputes about “what time really is” do not cash out as different predictions. They cash out as different stories wrapped around the same equations.

That does not make the question trivial. It means the question is underspecified. “Iteration” can mean at least four different things: (i) a relabeling of a time parameter, (ii) a fundamental discrete tick, (iii) a computational update rule with resource limits, or (iv) an agent-centered sequence of belief updates. Each implies a different physics-facing commitment. Likewise “time” can mean a coordinate in a model, an objective flow, a thermodynamic arrow, or an emergent bookkeeping variable tied to clocks.

This paper widens the frame. Instead of asking which of two labels is “true,” we ask what alternative order-structures could underwrite the appearance of change, memory, causation, and prediction. We then ask what would distinguish them operationally. The goal is not to pick a favorite metaphysics. The goal is to prevent category errors: mistaking a change of vocabulary for a change of physics, and mistaking a felt feature of experience for a necessary primitive of the world.

1.1 The definitional trap: re-labeling versus new physics

A large fraction of “time vs iteration” debate is a definitional trap: two descriptions can be different in language while identical in empirical content. If “iteration” is simply an index n that orders states in the same way time t does, and there exists a monotone mapping $t = f(n)$ that preserves all correlations among observables, then nothing in principle distinguishes them. The world supplies correlations; we supply a parameter that makes those correlations easy to write down. In that situation, arguing whether the parameter is “time” or “iteration” is like arguing whether a road trip is measured in miles or kilometers. The physics lives in what stays invariant under re-description.

This trap is easy to fall into because the words carry extra metaphysical baggage. “Time” sounds continuous, flowing, and universal. “Iteration” sounds discrete, computational, and stepwise. But those connotations do not automatically attach to the mathematics. A continuous parameter can describe a fundamentally discrete process as an approximation. A discrete index can describe a perfectly continuous dynamical system by sampling it. A “tick” can be a numerical scheme, not an ontological atom.

To escape the trap, each term must be sharpened into a claim with consequences. For example:

- If “iteration” means a fundamental smallest update step, then it had better imply some observable departure from exact continuity at some regime, even if extremely high-energy or extremely small-scale.
- If “iteration” means a literal computational rule, then it should entail constraints or artifacts that are not already implied by standard physical laws.
- If “time” means objective flow (a “moving now”), then it should add structure beyond a dimension in a spacetime model and beyond mere ordering; it should have signatures that cannot be removed by reparameterization.

Without such sharpening, the debate is not wrong, but it becomes under-determined: many incompatible ontologies fit the same observable data.

1.2 What physics can and cannot decide

Physics can decide questions that create empirical inequivalence. If two ontologies produce different predictions for measurable outcomes, physics can in principle choose between them (or constrain them) by experiment and observation. Physics can also decide questions of internal consistency: whether a proposed framework can reproduce known symmetries, conservation laws, and robust features of established theories.

But physics cannot decide between two descriptions that are empirically equivalent and mathematically intertranslatable. In such cases, the choice becomes pragmatic (which description is simpler, more compositional, more explanatory) or philosophical (which ontology one prefers). That boundary matters here because several prominent approaches treat “time” as a parameter-choice issue: you can write the same relational content using different clocks, or different parameterizations, without changing any predictions.

So the key is to separate three layers:

- Operational layer: what is measured (correlations among observables, clock readings, causal influence, record statistics).
- Theoretical layer: the mathematical structure that generates those measured correlations (dynamical laws, constraint equations, symmetry principles).
- Ontological layer: the story about what is “really” fundamental (time, causality, records, computation, emergence).

Physics is strongest at the first two layers. It can often constrain the third layer only when the ontological commitment forces distinctive structure into the first two.

This is why “time vs iteration” is too narrow: it invites an ontological verdict before the operational discriminators are specified. A better method is to enumerate candidate primitives for ordering and then ask: what must be true if that primitive is fundamental, and what could we see that would differ if it is not?

1.3 The paper’s aim: a taxonomy of orderings of change

The aim of this paper is to map a taxonomy of ways the world could exhibit ordered change without presupposing that “time” or “iteration” is the primitive. The taxonomy is built around a simple question:

What is the minimal structure required to make sense of “before/after,” “cause/effect,” “memory/record,” and “prediction,” and what alternative primitives could supply that structure?

This framing yields multiple candidates for “ordering” that are not reducible to the time/iteration binary:

- Correlation-first orderings, where clocks are subsystems and “time” is a relational mapping.
- Causal partial orders, where “influence” is primitive and global simultaneity may be absent.
- Thermodynamic/entropic orderings, where the arrow arises from boundary conditions and the growth of stable records.
- Record-based orderings, where what exists fundamentally are present constraints among traces that encode apparent histories.
- Gauge/reparameterization views, where time-parameters are not physical until fixed by a choice of clock.
- Phase and cyclic orderings, where an angle-like variable orders return and recurrence.
- Scale orderings, where “before/after” is replaced by flow from micro to macro descriptions.
- Inference and computational-depth orderings, where “later” can mean “harder to derive,” “more structured,” or “more compressed.”
- Agency/decision orderings, where sequence is defined by information availability and control-loop updates.

The paper's contribution is not a new speculative theory, but a disciplined map: a way to avoid speaking past each other, and a way to locate "time" and "iteration" as special cases inside a larger space of order-structures. We will keep returning to a practical test: whenever a proposal claims fundamental status, what observable or structural commitments follow, and what would count as evidence against it?

In short: the paper is a vocabulary and a filter. It widens the ontology menu, but tightens the standards for claiming that any item on the menu is physics rather than poetry.

2. What Counts as an Ordering in Physics

Before we can compare "time," "iteration," causality, entropy, records, phase, or scale as candidates for fundamental ordering, we need a criterion for what "ordering" even means in physics. Physics does not begin from metaphysical labels; it begins from measurement and from the stable regularities that survive changes of description. An "ordering" is physically meaningful only insofar as it (i) organizes observable events into a consistent structure, (ii) supports counterfactual reasoning (what would happen if conditions differed), and (iii) remains robust under transformations that should not change the physics (coordinate changes, gauge choices, reparameterizations, relabelings).

This section builds the minimal toolkit. It separates parameters from observables, highlights the central role of invariants, and explains why global order is often too strong a demand. It then treats clocks not as abstract external timekeepers but as physical subsystems embedded in the same world they measure—subject to noise, drift, and environment. Finally, it offers an operational definition: in physics, "order" is what we can extract from measurable correlation structure, not what we can name.

2.1 Parameters, observables, and invariants

A common source of confusion in time debates is mistaking a parameter for a physical observable. A parameter is a piece of mathematical scaffolding used to label points along a curve, or to index solutions of equations. An observable is something that can, in principle, be measured by an apparatus and reported as a number (with uncertainties) in an experimental context. Between these lies a third category that often matters most: invariants—quantities or relations that remain unchanged under the permissible transformations of the theory.

In classical mechanics, "time" is often introduced as an independent variable t that indexes the system's evolution $x(t)$. But the physics is not t itself; the physics is in the relations among measurable quantities: positions, momenta, energies, trajectories, scattering angles, transition rates, and so on. You can often reparameterize the same physical trajectory with a different

parameter s and obtain the same set of observable correlations, provided s is monotone and the dynamical laws are reexpressed accordingly. In such cases, the parameter is not physical. It is a label.

This distinction becomes sharper in relativistic settings. Special relativity denies a universal global time; different observers slice spacetime differently. Yet physics remains coherent because it is anchored in invariants: proper time along a worldline, invariant intervals, conserved quantities, and frame-independent predictions for experimental outcomes. The lesson is not “time is fake,” but “time is not a single thing.” Some time-like quantities are coordinate conventions; others are operationally defined intervals measured by physical clocks along specific paths.

The correct question for any proposed ordering primitive is therefore: what is invariant? If “iteration” is fundamental, what invariant structures does it impose that cannot be transformed away? If “causal order” is fundamental, what invariant relation captures it (e.g., reachability/influence) and how does it map to what instruments measure? If “entropy arrow” is fundamental, what invariant directionality does it supply and under what conditions does it fail?

An ordering becomes physics when it is not merely a convenient index but a constraint on possible correlations—something that survives changes of coordinate, notation, or description.

2.2 Global order, local order, and partial order

A naive picture of ordering assumes a single total sequence: every event is either earlier or later than every other event. That is a global linear order. It is psychologically natural, but physically too strong. Relativity already teaches that not all pairs of events admit an observer-independent “before/after.” Events can be spacelike separated, meaning different observers can disagree on their temporal order without contradiction. Physics survives because it does not require a global total order. It requires consistent local structure.

This leads to a more general notion: partial order. In a partial order, some pairs are comparable (“A precedes B”), while others are incomparable (neither precedes the other in a meaningful, observer-independent way). Causal structure is naturally expressed this way: if A can influence B (signals, causal propagation), then A is in the causal past of B. If no influence relation exists (spacelike separation), the events are not causally ordered. That does not mean “nothing happened”; it means the ordering relation is not defined for that pair.

The “global vs local” distinction also matters in statistical and thermodynamic contexts. The arrow of time is not a global axiom stamped everywhere equally. It is a pervasive regularity that emerges from boundary conditions and from the behavior of large numbers of degrees of freedom. Locally, small systems can fluctuate. Microscopically, dynamics may be time-reversal

symmetric while macroscopic irreversibility emerges as a coarse-grained fact. If one treats the arrow as an ordering relation, it may be probabilistic and context-dependent rather than absolute.

So, when someone says “the universe is an iteration,” we should ask: iteration of what, and with what scope? Is it a global tick affecting all events simultaneously (a total order)? Or is it local updating of regions (a partial order)? If updates are local, then “iteration” begins to resemble a causal network more than a universal counter. Conversely, when someone says “the universe is in time,” we should ask: do they mean a global time parameter (ruled out in standard relativistic understanding), or do they mean locally measurable durations (proper times) along physical processes?

A mature notion of ordering in physics is therefore plural: it allows local orders, partial orders, and multiple incompatible but consistent slicings of the same underlying structure.

2.3 The role of clocks as physical subsystems

A clock is not a metaphysical primitive. A clock is a physical system whose internal state changes in a regular way and can be coupled to a readout. Pendulums, atomic transitions, oscillators, decay processes, orbital periods, and even pulsars can serve as clocks in different regimes. The key point is that clocks are themselves part of the universe; they are not outside it providing an external parameter.

This matters for two reasons.

First, it forces the “time” question to become concrete. When we speak of time operationally, we mean comparisons of physical processes: “System S changes by amount ΔS while clock C changes by amount ΔC .” Time is the correlation between changes in S and changes in C. In this relational view, different clocks can disagree (due to environment, relativistic effects, noise, or coupling), and the job of physics is to predict and reconcile those disagreements within a consistent framework.

Second, it reveals what “ordering” must provide: it must explain why stable clocks exist and why their readings can be used as a reliable coordinate for other processes. If an ontology denies fundamental time, it still must account for the emergence of clock-like subsystems and for the near-universality of timekeeping in our experience. If an ontology posits discrete iteration, it must explain why many clocks appear continuous and why the apparent continuity survives at scales far above the hypothetical tick.

Treating clocks as subsystems also highlights an important subtlety: you do not measure “time” directly. You measure a clock variable and interpret it as time under a model. The robustness of timekeeping is therefore a mixture of physics and convention: physics constrains how good

clocks can be, how they behave under motion and gravity, and how they drift; convention defines standards and synchronization protocols.

A practical way to frame it: if you change your clock choice, what remains invariant? The answer to that question tells you what ordering is truly physical and what is representational.

2.4 Operationalism: “order” as measurable correlation structure

Operationalism here does not mean “only measurement is real” in a philosophical sense. It means something more modest and more powerful: in physics, if you cannot say what would be measured differently under alternative hypotheses, you do not yet have a physics claim. You have a semantic preference.

So what is “order,” operationally? It is the structure of correlations among observables that allows consistent reconstruction of “before/after,” “duration,” “causal influence,” and “record consistency.” Different ordering primitives emphasize different measurable structures:

- A time-parameter picture emphasizes trajectories and rates: how observables change as a function of a chosen clock variable.
- A causal-order picture emphasizes influence and signaling constraints: which events can affect which other events.
- A thermodynamic-arrow picture emphasizes statistical asymmetries: the direction in which coarse-grained entropy increases and records accumulate.
- A records picture emphasizes mutual consistency among localized traces: the present contains correlated “imprints” that behave as if they refer to a shared past.
- A gauge/reparameterization picture emphasizes invariance: different parameter choices yield the same set of measurable correlations.
- A phase-order picture emphasizes periodic correlations: stable cycles that provide ordering without linear progression.
- A scale-order picture emphasizes universality across description levels: how behavior changes as you move from micro to macro coarse-grainings.
- An agency-order picture emphasizes information availability and control: how an agent’s prediction/decision sequence structures experience.

By this operational standard, an ordering “counts” in physics if it does at least one of the following:

1. It constrains which correlations are possible.

2. It predicts new correlations not predicted by rival orderings.
3. It explains why certain correlations are stable and portable (e.g., why clocks work, why records persist).
4. It provides invariant quantities that remain fixed under legitimate transformations of description.

This gives us a disciplined way to proceed. Rather than arguing whether “time” or “iteration” is the deeper word, we ask: what correlation structures must exist if a given ordering is fundamental, and what correlation structures would fail (or change) if it is not?

That is the physics-facing core of the paper: order is not primarily a metaphysical noun. It is a measurable pattern in how observables co-vary—and the invariants that survive our many ways of labeling them.

3. Time as a Coordinate Versus Time as Becoming

A central ambiguity in everyday speech is that we use the single word "time" to name two very different things. First, time can mean a coordinate-like parameter used to organize change: a way to label events, compare durations, and write laws in a compact form. Second, time can mean becoming: the felt sense that reality "moves" from future to present to past, as if the present is objectively special and advances.

Physics handles the first meaning extremely well. The second meaning is harder: it is not obviously a physical variable, and it is not automatically implied by the success of time-coordinate formalisms. This section separates these notions cleanly, identifies where the conceptual pressure points sit, and clarifies what would have to be true for "flow" to be something physics could treat as real structure rather than experiential interpretation.

3.1 The block-universe stance: time without flow

The block-universe stance treats the world as a complete set of events related by structural constraints (geometry, causality, dynamical laws), without any additional global "moving present." On this view, time is part of the description of the world, not a process that happens to the world.

What is gained by this stance is conceptual economy. You do not need a second time (a meta-time) to explain how time itself "passes." You do not need an objective "Now" stamped across the universe. You only need:

- A set of events (or states) that can be related.

- Rules that constrain which event-patterns are allowed.
- Local notions of duration along physical processes (for example, what a clock measures along its own history).

In a block picture, it is natural to speak tenselessly: the event of a supernova, the event of your breakfast, and the event of tomorrow's sunrise are all elements of the block. They differ in where they sit relative to one another, not in whether they "exist yet." The language of "yet" is treated as an artifact of perspective.

This does not force a single global ordering of everything. Modern physics is comfortable with local durations (what particular clocks measure) and with observer-dependent slicings. Two observers can disagree about which distant events are "simultaneous" without any contradiction, because simultaneity is not taken to be a global physical primitive. What remains stable are invariant relationships: causal connectibility, local clock intervals, and the reproducible correlations that experiments reveal.

A helpful way to see the block stance is: the world is the total constraint-satisfying structure; "dynamics" is one way of describing relationships inside that structure. The equations can often be written as evolution rules ("given this, predict that"), but the block stance interprets this as a method for extracting conditional relations, not as proof of an objective universal flow.

3.2 Becoming/flow as an additional claim

When someone insists that time is not merely a coordinate but "really passes," they are making an additional claim beyond what the block stance needs. The key is to state that additional claim precisely. There are several distinct versions:

- A distinguished present: there is an objective fact of the matter about what is "Now," not merely "now for me given my clock and light-cone limitations."
- Objective passage: the distinguished present changes, and the change is not just in our knowledge or perspective.
- Ontological asymmetry: the past and future differ in kind, not only in our access to records.

Each of these adds structure. In particular, objective passage seems to demand some parameter with respect to which the present advances. But if that parameter is itself time-like, we risk an infinite regress: what makes that meta-time pass? If it is not time-like, we must specify what it is and how it couples to physical observables.

This is why "flow" is not automatically a physics claim. If you cannot say what additional measurable structure comes with it, you have not yet moved from metaphysics to physics.

Becoming can be framed in less ambitious ways that avoid regress:

- Growing-block views: the past is fixed, the present is the edge, and the future is not yet part of reality. This still requires a rule for what becomes fixed and when.
- Local becoming: only along certain worldlines (agents, clocks) does "now" have meaning; the universe has no global now, only many local nows. This tends to collapse back toward a relational view unless it also claims new observables.
- Process metaphysics: the primitives are processes rather than events. But to count as physics, "process" must still be expressible as constraints on correlations among observables.

So the paper will treat "becoming" as a family of proposals that share a common burden: they must specify what is added to the coordinate-time picture, and they must explain why that addition does not simply rename familiar features (entropy increase, memory asymmetry, or local clock readings).

3.3 The "experience gap": why flow feels unavoidable

If physics can do so much with time-as-coordinate, why does flow feel so unavoidable?

Because the brain is not a passive camera. It is a prediction-and-control system embedded in an environment with strong asymmetries. Several features conspire to create the phenomenology of passage:

- Memory asymmetry: we carry accessible records of what we call "past" states (both internal memory and external traces), but not of "future" states. This is not merely psychological; it is an environmental regularity tied to record formation and stability.
- Prediction asymmetry: we can reliably forecast some patterns forward (sunrise, inertia, periodicity) but cannot "retrodict" in the same way because retrodiction is constrained by present records, not by forward evolution from unknown initial microstates.
- Updating: an agent experiences a sequence of belief updates. Even if the world were a block, an agent's internal state can be ordered by its own successive revisions. The felt "now" is closely tied to the boundary between "incorporated into the model" and "not yet incorporated."
- The specious present: conscious experience integrates across a short temporal window. We do not experience instantaneous points; we experience a smear of recent data fused into a coherent present. That fused window makes the present feel like a platform from which the world is continually arriving.

- Irreversibility in practice: macroscopic processes generate heat, dissipate energy, and create stable records. This makes the forward direction special for agents: it is the direction in which actions propagate into the world and get written into durable traces.

These points narrow the gap. The sense of flow can be understood as a feature of how bounded agents navigate an entropic environment: we are information-limited systems that maintain models, update them, and act to reduce surprise and maintain viability. In that setting, "passage" is the agent-level signature of continual model revision coupled to irreversible record formation.

But there is still a residual gap: even if you explain memory and irreversibility, many people feel that you have not captured the "metaphysical push" of time, the insistence that the present is objectively special, not merely psychologically central. That insistence is exactly the point where becoming goes beyond what coordinate-time requires, and where it must either (i) become a measurable structure, or (ii) remain an interpretation layered on top of empirically adequate physics.

3.4 What would it mean for flow to be physical

For "flow" to be physical, it must do more than mirror the agent's update stream or the thermodynamic arrow. It must correspond to additional structure that, at least in principle, could be detected as a difference in the correlation structure of observables.

That pushes us toward a discriminator: a physical flow should imply at least one of the following.

1. A distinguished foliation or preferred frame

If there is a real global present, the universe must be sliced into "space at now" in a preferred way. That would be extra structure beyond the standard coordinate freedom. It might show up as:

- A measurable "rest frame" that is not merely conventional.
- Tiny direction-dependent effects in propagation or dynamics.
- A new field-like object that picks out the foliation.

The danger is that such structure often clashes with the deep symmetry idea that no inertial frame is fundamentally preferred. If "flow" demands a preferred slicing, it must explain why the usual symmetry works so well and why deviations have not already been seen.

2. Time-asymmetric fundamental dynamics

If the microscopic laws are not time-reversal symmetric in a deep way, that asymmetry

could ground an objective arrow that is not merely statistical. In principle, this could manifest as:

- Irreversibility that cannot be removed by changing description or coarse-graining.
- Violations of reversibility expectations even in carefully isolated systems.
- A fundamental "loss of information" term that cannot be interpreted as ignorance.

But this is a heavy claim. Many ordinary irreversibilities are compatible with reversible micro-laws plus coarse-graining. A physical flow would need to be an asymmetry that remains after you account for all standard sources of apparent irreversibility.

3. **A real "present-making" variable**

A true becoming theory might posit a variable P that indicates "presentness" and changes in a way that is not reducible to coordinate choice. For physics, this variable must couple to matter or measurement so that it affects outcomes. Otherwise it is epiphenomenal: it exists in the story but does no work.

The coupling is the key. If P never changes any observable, it cannot be empirically distinguished from its absence.

4. **A measurable limit on temporal divisibility (a tick)**

If flow is tied to discrete becoming, there may be a smallest meaningful update step. If this discreteness is fundamental (not an approximation), it could leave signatures like:

- Deviations from perfect continuity in high-precision timing experiments.
- Additional noise floors in phase-sensitive measurements.
- Small modifications to propagation at extreme energies or long baselines.

Again, the burden is not to assert these effects, but to specify them sharply enough that they become constraints rather than poetry.

5. **A difference between "future" and "past" beyond records**

In a purely record-based or thermodynamic picture, the difference between past and future is largely about accessible traces and boundary conditions. A physical becoming view might claim that the future is not merely unknown, but literally not fixed. To be physics, that literal openness must change something measurable: for example, it might imply objective indeterminism with a specific signature, or constraints on retrodiction that go beyond standard statistical limitations.

The paper's posture will be disciplined: "flow" is not dismissed, but it is not granted for free. If flow is real, it should either:

- Produce a preferred structure (frame, foliation, variable), or
- Produce a fundamental asymmetry in dynamics, or
- Produce a discrete update signature, or
- Produce a measurable distinction between past and future not explainable by records and boundary conditions.

If none of these emerge, the most conservative conclusion is that physics supports time-as-coordinate (and local clock-measured durations), while flow remains an experiential and interpretive layer. That is not an insult to experience. It is a separation of categories: physics describes invariants of correlation; experience describes the lived interface a bounded agent has with those correlations.

This sets up the next sections: we will compare proposals that treat correlation, causality, entropy, and records as the ordering primitives, and we will keep asking the same question: what additional structure do they posit, and what would it mean operationally for that structure to be there.

4. Iteration: Discreteness, Update Rules, and Hidden Assumptions

"Iteration" sounds clean: a world that updates step by step. But the word is a suitcase packed with hidden assumptions. Sometimes "iteration" means nothing more than a new label on an old parameter. Sometimes it means a fundamental discreteness, a smallest tick that cannot be subdivided. Sometimes it means a literal computational procedure, as if the universe is running an update rule with finite resources. And sometimes, without noticing, it smuggles in a preferred global synchronization—an absolute "frame" in which the updates occur simultaneously.

This section separates these senses and makes the commitments explicit. The aim is to keep iteration from being a vibe and to force it into the same standard we applied to time: if iteration is a physical primitive, what invariants, constraints, and observable signatures follow?

4.1 Iteration as mere reparameterization

The weakest and most common meaning of iteration is reparameterization: instead of writing states as $x(t)$, you write them as $x(n)$, where n is an index. If there is a monotone mapping $t =$

$f(n)$ (or $n = g(t)$) such that all observable correlations are preserved, then nothing has changed in the physics. You have only changed the labeling.

This is often what happens when people describe continuous dynamics “as if” it were iterative: you sample a continuous trajectory at discrete points, or you write a differential equation as a limit of a difference equation. That can be computationally useful, but it does not imply that the world itself updates in discrete steps. It implies only that a discrete approximation can track the world to whatever precision you need.

A telltale sign that iteration is being used as reparameterization is that it carries no new constraints. The theory still allows arbitrarily fine subdivisions of the parameter (or arbitrarily small time intervals), and there is no physical scale at which the “step” becomes meaningful. In this case:

- “Iteration” is descriptive scaffolding.
- The only physical content remains in invariants and correlations among observables.
- No experiment could distinguish “ t ” from “ n ” because the mapping is a matter of convention.

This version of iteration is harmless and often clarifying—especially in relational or gauge-like contexts where the parameter is not observable anyway. But it is not an ontological claim.

4.2 Iteration as fundamental discreteness (a true tick)

A stronger claim is that there exists a smallest physically meaningful update step: a true tick. The world does not merely admit discrete descriptions; it is discrete at base. Between tick k and tick $k+1$ there is no intermediate physical state.

This is a major ontological commitment because it introduces a new kind of scale: the tick size. Even if that scale is unimaginably small, it is a real structure that must either (i) respect the symmetries we observe extraordinarily well, or (ii) break or deform them in ways that are constrained.

A true tick can be interpreted in two importantly different ways:

- **Discrete state with continuous geometry:** the state updates discretely, but spacetime remains effectively continuous as an emergent or approximate description.
- **Discrete spacetime (lattice-like):** both state and “where/when” are discrete, potentially suggesting a combinatorial substrate.

Either way, a true tick threatens to leave fingerprints. The fingerprints can be subtle, but the basic logic is: if there is a smallest update, then at sufficiently high resolution or extreme conditions, the smoothness assumptions of standard theories cannot remain exact.

What kinds of signatures might a true tick imply (in principle)?

- **Departures from exact Lorentz symmetry:** if updates occur in synchronized steps, “simultaneous update” can privilege a slicing of spacetime.
- **Modified propagation at extreme energies:** discrete updating can produce tiny deviations from smooth wave propagation that accumulate over long distances.
- **Additional noise floors:** if evolution proceeds in discrete jumps, there may be irreducible jitter in phase- or timing-sensitive experiments.
- **Quantization artifacts beyond known quantization:** this is not ordinary quantum discreteness of measurement outcomes; it is discreteness of evolution itself.

A true tick is not ruled out merely because everyday physics looks continuous. Continuity can emerge from discreteness the way fluid flow emerges from molecules. But the burden shifts: a discrete ontology must explain why continuity works so well and why symmetry violations (if any) are so suppressed.

4.3 Iteration as computation: update rules and resource bounds

An even stronger meaning says: the universe is not only discrete, it is computational. There is an update rule U that takes the state at step n to the state at step $n+1$. This is tempting because it imports a familiar explanatory template: programs generate behavior.

But “computation” is not a single claim; it is a bundle. To say the universe computes is to risk importing assumptions from our machines that may not apply. The useful way forward is to separate three layers:

- **Rule-like evolution:** there exists a lawful mapping from one state to the next.
- **Finite resources:** the mapping is executed with bounded memory, bounded information flow, bounded precision, or bounded bandwidth per step.
- **Algorithmic constraints:** certain kinds of global optimization or “instant” coordination are impossible because the update is local and bounded.

The first layer—rule-like evolution—is basically what physics already does: laws map conditions to constraints on outcomes. The second and third layers are where “computational iteration” becomes a distinct physical claim. A resource-bounded universe might imply:

- **Maximum information throughput:** constraints on how much state can influence how much other state per step (related to locality and finite signal speed, but potentially stricter).
- **Precision limits:** an irreducible finite resolution in state variables, not merely practical measurement limits.
- **Cost of coordination:** global correlations cannot appear “for free”; they require time/steps/communication, making certain patterns impossible without sufficient depth.
- **Complexity-style arrows:** the direction of “later” might align with increasing computational depth, not merely increasing entropy.

However, computation talk also hides a trap: any dynamical system can be described “as if” it were computing. That does not mean it is computing in a resource-bounded, algorithmically constrained sense. To make computation physical, you must specify what resource is limited, how it is counted, and what predictions change because of the limitation.

In short: “the universe has laws” is not “the universe runs an algorithm,” unless “algorithm” implies additional constraints beyond ordinary locality and conservation.

4.4 Where iteration hides a preferred frame or structure

The most common hidden assumption in iteration talk is global synchronization. People imagine a universal update: the entire universe advances from step n to step $n+1$ together. That is an implicit preferred slicing of reality into global “instants.” Once you posit that, you have smuggled in a structure that standard relativistic physics works hard to avoid: a universal present defined everywhere at once.

There are ways to avoid this, but each changes what “iteration” means:

- **Local iteration:** regions update based on local conditions, not in perfect global lockstep. This begins to resemble a causal network or event graph rather than a global counter.
- **Asynchronous iteration:** different parts update at different times with partial ordering; the fundamental structure is then closer to causality than to a universal tick.
- **Emergent synchronization:** global “steps” are an approximation that becomes meaningful only in certain regimes (for example, when a chosen frame is convenient for computation or cosmology).

If one insists on a globally synchronized tick, one must confront the question: synchronized relative to what physical standard? If the answer is “relative to itself,” that is not yet an operational definition. It becomes operational only if the tick defines measurable rest,

anisotropy, or propagation effects. If those effects are absent, the “global tick” risks being an unobservable structure—again, story without consequence.

A second hidden structure is **a lattice or grid**. Many people visualize discrete iteration as occurring on discrete space. But a grid tends to introduce directionality (axes), anisotropy, or preferred lengths unless carefully constructed. If the world’s discreteness is grid-like, why do we observe such smooth isotropy and symmetry? The discrete theory must either produce symmetry as emergent to extremely high accuracy or replace exact symmetry with a deformed symmetry that matches observations.

So the warning is simple: iteration is not a single alternative to time. It is a family of claims, and the most intuitive versions often import a preferred “now” by stealth. If iteration is to compete with time-as-coordinate in a physics-respecting way, it must either:

- be a mere reparameterization (and then it is not a competing ontology), or
- be discrete but symmetry-respecting (and then it owes an emergence story), or
- be computational with explicit resource bounds (and then it owes testable constraints), or
- be local/asynchronous (and then it collapses toward causal partial order rather than global ticking).

This sets the stage for the next sections: relational time and causal order already provide rigorous notions of ordering without global synchronization. The key question is whether “iteration” adds anything over and above those structures—or whether, once purified of hidden assumptions, it simply names one region of the larger taxonomy of orderings of change.

5. Correlation-First Ordering (Relational Time)

Correlation-first ordering begins from a simple refusal: do not assume a universal external time parameter as a primitive. Instead, treat “time” as something extracted from correlations among physical degrees of freedom. In this stance, the world is made of interacting subsystems. Some subsystems can function as clocks because they exhibit regular, monotone, and robust change. Other subsystems are “systems of interest.” What we call time is the mapping that aligns changes in the system of interest to changes in the chosen clock subsystem.

This approach is not merely philosophical minimalism. It matches how measurement actually works. No experiment measures “t” directly. Experiments compare one process to another: the phase of an oscillator, the count of atomic transitions, the ticks of a device, the arrival times of pulses. The relational stance says: take that seriously. If time is always read off from clocks, then

clocks should be in the theory as physical objects, and “time” should be defined as a relation, not as an external stage on which the play occurs.

Relational time is also a natural response to the underdetermination we saw earlier. If “time vs iteration” collapses into reparameterization, then what remains physically meaningful are correlations. Correlation-first ordering elevates that remainder to the primitive.

5.1 “Time from clocks”: internal time variables

In a correlation-first picture, a clock is a subsystem with a state variable C that changes in a way we can track and use as an index. The system of interest has variables S . Instead of writing $S(t)$, we write $S(C)$: the state of S as a function of the clock reading C .

Operationally this means: you pick a physical process to serve as clock, and you define time by that process. Conceptually it means: “time” is not something the universe has; it is something agents construct by using one part of the universe to parameterize another part.

A minimal internal-time construction has three ingredients:

- **A clock variable that is monotone over the regime of interest.** If the clock variable cycles or stalls, it cannot globally order change, though it may still be useful locally (phase clocks are a special case we treat later).
- **A coupling that allows the clock reading to be compared with the system state.** If the clock is completely decoupled, it is useless for describing the system.
- **A stability condition:** small perturbations should not destroy the clock’s usefulness, or at least we must know how the clock drifts.

In everyday physics, this is almost too obvious: we use cesium transitions, quartz oscillators, mechanical resonators, pulsars, planetary motion, or decay counts as clocks. Relational time simply promotes this from a practical habit to an organizing principle.

A key payoff is that relational time makes explicit what is often hidden: any talk about “the system evolving in time” is shorthand for “the system’s state correlates in a regular way with some clock’s state.” The world provides correlations; “time” is the coordinate we build from them.

5.2 Multiple clocks and the problem of clock choice

The moment you take clocks seriously as physical subsystems, you inherit the clock-choice problem: there is not one perfect clock; there are many imperfect clocks. Different clocks can disagree, drift, or respond differently to environment, motion, and gravitational conditions.

Even in idealized theory, different internal-time choices can lead to different descriptions of “evolution.”

This does not mean the relational program fails; it means the notion of time becomes plural and contextual.

There are several distinct sources of clock multiplicity:

- **Practical multiplicity:** different clock technologies have different noise, drift, and sensitivity to conditions.
- **Relativistic multiplicity:** different worldlines accumulate different clock readings due to motion and gravity; there is no universal synchronization that all clocks share.
- **Model multiplicity:** different choices of internal variable as clock can simplify different aspects of the dynamics.

The core conceptual issue is: if time is “whatever the clock says,” what makes some clock choices better than others? Physics answers this in a disciplined way: a good clock is one whose correlations with other subsystems are stable, whose behavior is well-modeled, and whose readings can be transported and compared with controlled uncertainty.

But relational time introduces a deeper point: even “perfect” clocks are perfect only relative to a theory and a regime. A clock is not a Platonic axis. It is a dynamical subsystem. So clock-choice is not merely instrumental; it is part of what any time concept must account for.

Relational frameworks often handle this by privileging invariants: when you change clock variables, the observable predictions should remain consistent, even if the intermediate descriptions differ. This is the same moral as gauge freedom: different choices of parameterization are allowed, but the physics is what survives those choices.

5.3 When relational time mimics ordinary time perfectly

Relational time can look indistinguishable from ordinary time whenever two conditions hold:

1. There exists at least one clock variable that is sufficiently monotone and regular over the regime considered.
2. The dynamics of the system relative to that clock reproduces the same observable predictions as standard time-parameterized physics.

In ordinary laboratory practice, these conditions are routinely satisfied. A high-quality atomic clock behaves so regularly that we treat its reading as “time” and write other processes as functions of that reading. The relational interpretation changes nothing operationally; it changes only what we claim is fundamental.

This is an important point: relational time can be empirically equivalent to coordinate-time physics. It can “mimic” ordinary time so well that no experiment distinguishes them, because both frameworks predict the same correlations among all measurable quantities. In that case, the debate is not about predictions. It is about metaphysical bookkeeping.

However, relational time becomes more than reinterpretation in regimes where ordinary time becomes ambiguous or problematic:

- **When no single global time parameter is available or natural** (as in strongly relativistic contexts).
- **When the relevant degrees of freedom include the clock itself** (the clock cannot be treated as external).
- **When the fundamental description is constraint-based rather than evolution-based** (where “change” is expressed as relations among variables rather than explicit time progression).

Even without naming specific formalisms, the conceptual pivot is clear: if you cannot assume a background time, but you can still construct stable internal clocks and express predictions as correlations, then you can do physics without fundamental time.

And if you can do physics without fundamental time while retaining the same empirical content, then “time as fundamental” becomes a choice, not a necessity.

5.4 Conceptual payoffs and costs: realism about parameters

Relational time offers a set of payoffs that align with the discipline of invariants, but it also imposes conceptual costs.

Payoffs:

- **No surplus structure:** you do not posit an invisible universal time axis if all measurements reduce to comparisons among physical processes.
- **Clocks become first-class:** the theory aligns with practice by treating timekeeping as physics, not as a backdrop.
- **Natural fit with local/partial ordering:** it meshes cleanly with relativity’s refusal of global simultaneity and with causal partial orders.
- **Clarity about gauge-like redundancy:** it makes it easier to see when disputes are just reparameterizations.

Costs and tensions:

- **Plurality of times:** if different clocks can define different “times,” you must explain why a shared time standard is possible and why it works so well in practice.
- **The meaning of “evolution”:** if time is internal and chosen, “evolution” is not a primitive story; it is a derived description relative to a selected clock. Some intuitions about universal progression must be surrendered.
- **The experience gap remains:** relational time explains measurement structure, but it does not automatically explain the felt flow of becoming. It can accommodate that phenomenology, but it does not force it.
- **Risk of over-deflation:** some readers may feel relational time “explains away” time. The more careful statement is that it relocates time: from a global stage to a relational coordinate built from physical processes.

The realism question is: are parameters real, or are they just tools? Relational time leans toward a disciplined realism about observables and invariants, and a disciplined anti-realism about bare parameters. It does not deny that time talk is legitimate; it denies that the bare symbol used to label change is automatically an ontological primitive.

This matters for our larger taxonomy. Once you accept correlation-first ordering, “iteration” becomes easier to classify: iteration is either (i) another parameter choice (a clock index), or (ii) a claim about discreteness or computation that must show up as altered correlations. And “time-as-flow” becomes easier to isolate: flow must be additional structure beyond correlations among clock readings and system behavior.

So relational time functions as a conceptual clarifier: it strips ordering down to what experiments actually give us—correlations—and forces every stronger claim (global time, true ticks, objective flow) to earn its keep by specifying what it adds that cannot be transformed away.

6. Causal Order as the Primitive

Correlation-first ordering says: start from measurable correlations and treat “time” as a convenient coordinate built from clocks. Causal-order-first goes one step more structural: it says the primitive ordering relation is not “clock reading” but “influence.” Instead of asking which event happens first in a global sense, it asks which event can affect which other event. The basic data of the world, on this view, is a directed relation: A can be in the causal past of B.

This is an attractive primitive because it matches what seems most physically non-negotiable: there are constraints on signaling and influence. Even if observers disagree about simultaneity, they do not disagree about whether influence is possible between two events (given the same underlying theory). Causal order also fits the intuition that “before/after” is meaningful chiefly where causal connection exists, and ambiguous where it does not.

The central promise of causal primacy is bold: perhaps spacetime geometry (distances, durations, even the notion of a “time coordinate”) is not fundamental. Perhaps it is reconstructed from a deeper causal web.

6.1 Partial orders: influence without global simultaneity

A causal ordering is naturally a partial order. That is not a technical flourish; it is the core feature that makes the view compatible with relativistic constraints.

- If event A can influence event B, then A precedes B in the causal order.
- If A and B cannot influence each other (no causal connection), then they are incomparable: neither precedes the other in the causal sense.

This handles the absence of global simultaneity cleanly. “Simultaneous” becomes, at best, a coordinate-dependent convenience, not a universal fact. What matters is whether events can be connected by possible influence.

Two clarifications matter here.

First, “influence” does not mean human agency or intention. It means physical possibility: could a signal, interaction, or causal chain connect A to B under the constraints of the world? This is a structural relation among events.

Second, partial order does not mean disorder. Partial orders can be richly structured. They can encode layered causal cones, local neighborhoods, and nested regions of mutual influence. They can generate robust notions of “earlier” and “later” along connected chains without requiring that every pair of events be comparable.

If you adopt causal order as primitive, you stop trying to force the world into a single universal sequence. You accept that “before/after” is fundamentally local-to-relational: defined where influence makes it meaningful.

6.2 Reconstructing geometry from causality

If causal order is primitive, what becomes of geometry? The striking idea is that much of what we call spacetime structure may be recoverable from causal relations plus some notion of “counting” or “density” of events.

The intuition is simple:

- Causal structure tells you which events can connect to which (the skeleton of “cones” of influence).
- If you also know something like how many events occur in a region, you can infer something like volume, and from that infer aspects of distance and duration.

This is not the same as measuring distance with a ruler. It is inferring geometry as an emergent description of a deeper relational pattern.

Why might this be plausible? Because geometry in physics is not merely a picture; it encodes constraints on propagation and influence. Light cones, causal separation, and the structure of possible signaling are tightly linked to the shape of spacetime. So if you have the causal web, you already have a large part of what geometry does.

A causal-first stance therefore shifts the explanatory direction:

- Standard view: geometry is fundamental; causality follows from geometry.
- Causal-first view: causal relations are fundamental; geometry is a derived effective summary.

Operationally, this aligns with a key point from earlier sections: we often learn “geometry” from how signals propagate (light travel times, radar ranging, causal delays). Those are causal phenomena. So a causal-first ontology asks: why not treat the causal constraints as basic and let geometry be the convenient encoding?

The cost is that you must show reconstruction is sufficiently unique. Many different geometries can, in principle, share similar causal patterns unless you add additional structure (like event density or dynamics). So causal primacy typically comes paired with a second ingredient: a way to assign scales, measures, or typical separations so that geometry is not underdetermined.

6.3 Locality, signaling, and causal horizons

Causal order becomes physically meaningful when it is tied to locality and signaling constraints. “Locality” here means that influences propagate through chains of intermediate events; “signaling constraints” mean that there are limits on how rapidly influence can spread; “horizons” mean that some regions are causally inaccessible from others, at least within certain domains.

Three features matter.

- **Locality as compositional structure:** if causality is primitive, then “nearby” is defined by being connected through short causal chains, and “far” by requiring long chains. Locality becomes a property of the graph of influence, not a pre-given metric.
- **Finite influence speed as an ordering constraint:** the world can have causal boundaries not because we choose coordinates but because influence genuinely cannot cross certain separations without intermediate steps.
- **Horizons as structural facts:** if there are boundaries beyond which influence cannot pass (or cannot pass within certain regimes), then those boundaries carve the causal graph into regions with distinct accessibility properties. That has direct consequences for what can be known, controlled, or recorded.

From the standpoint of ordering, causal horizons are important because they define where ordering relations are meaningful. Inside a horizon, events can be compared by influence chains. Across a horizon, comparability may be limited or absent. This makes the causal view naturally aligned with the lived fact that knowledge and action are constrained: not everything is accessible; not everything can be synchronized; not everything can be causally “brought together” into a single narrative without assumptions.

In the taxonomy of orderings, causal order is a candidate for “why ordering exists at all.” It says: the minimal ordering is not about clocks; it is about permitted influence. Clocks then become particular subsystems that exploit stable local causal regularities to generate regular tick-like correlations.

6.4 How causal order differs from discrete iteration

Causal order is often confused with iteration because both can be represented as stepwise structures. But they differ in what they treat as fundamental, and in what they smuggle in.

Discrete iteration, in its naive form, suggests a global tick: a universal step $n \rightarrow n+1$ applied everywhere at once. That implies a global simultaneity structure. Causal order does not. It allows asynchronous, local advancement where “next” is defined only along causal links, not globally.

Here are the sharp contrasts.

- **Global vs local:** iteration typically presumes a global sequence of steps; causal order presumes only local precedence along influence chains.
- **Total vs partial:** iteration encourages total ordering (every event assigned to some step); causal order is naturally partial (some events incomparable).

- **Synchronization:** iteration often assumes synchronization; causal order treats synchronization as something to be constructed (if possible) from within the causal web.
- **Preferred frame risk:** global iteration risks introducing a preferred frame or slicing; causal order can remain compatible with the absence of a universal frame because it does not require global “now.”
- **Meaning of “step”:** in iteration, the step is a fundamental unit of update; in causal order, “steps” are just edges in an influence graph and can be heterogeneous (not all causal links are equal in “duration” unless additional scale structure is supplied).

This matters for the paper’s overarching argument. If one wants discreteness without preferred global time, causal-order-first is often the cleaner route: it permits discrete structure (events and links) without requiring global ticking. If one wants a computational update picture without global synchronization, causal order again provides the scaffolding: local update rules on a causal graph can be defined without a universal clock.

So causal primacy can be seen as a disciplined alternative to naive iteration: it keeps the intuition that “the world is built from events connected by influence,” while avoiding the hidden assumption that there exists a universal step counter that all events share.

In the next sections, we will compare causal order with entropy- and record-based orderings. The guiding question remains the same: what primitive relation is being posited, what does it buy, and what does it force you to give up—especially the temptation to reintroduce a global time by stealth.

7. Entropy and the Thermodynamic Arrow

If “time” as a coordinate orders events, and “causal order” links events by influence, the thermodynamic arrow addresses a different question: why does ordering feel *directed*? Why does the world present a robust asymmetry between “past” and “future” in practice—an asymmetry tied to memory, decay, dissipation, aging, and the one-way character of many macroscopic processes?

The thermodynamic arrow proposes that directionality is not a primitive flow of time, but a statistical regularity grounded in entropy: systems evolve from special, low-entropy conditions toward more typical, higher-entropy conditions under ordinary dynamics. This framework is powerful because it can explain why “earlier/later” correlates with the accumulation of records and the formation of memories without requiring a metaphysical moving present. But it also creates a burden: it must explain why the universe ever had low entropy in the first place, and it

must distinguish genuine irreversibility from apparent irreversibility due to coarse-graining and ignorance.

This section clarifies how entropy yields a direction, why records and prediction align with it, what “irreversibility” really means in physics terms, and why an arrow is not the same thing as a time parameter.

7.1 Low-entropy boundary conditions and “direction”

Entropy is often introduced as a measure of disorder, but for our purposes it is better understood as a measure of multiplicity: how many microscopic configurations correspond to the same macroscopic description. Low entropy means the macroscopic state is highly special: it corresponds to relatively few microstates. High entropy means the macrostate is typical: it corresponds to many microstates.

The crucial point is that the fundamental microscopic rules in many physical theories are compatible with reversibility: if you know the exact microstate, you can in principle run the dynamics backward. Yet macroscopic processes overwhelmingly exhibit a direction: ice melts, gases mix, heat flows from hot to cold, broken cups do not spontaneously reassemble. The thermodynamic arrow explains this direction as a consequence of probability plus a boundary condition:

- Given a low-entropy starting macro-condition and typical microstates compatible with it, evolution will almost certainly move toward macrostates with larger multiplicity (higher entropy).
- The reverse—spontaneous evolution from a typical high-entropy macrostate into a special low-entropy macrostate—is not forbidden by the micro-laws, but it is overwhelmingly improbable in practice.

The phrase “low-entropy boundary condition” is doing enormous work. It says: the universe (or the relevant closed system) began in a special macro-condition. That specialness supplies the arrow. If you had instead a boundary condition of maximal entropy, you would not get the familiar arrow; you would get equilibrium-like behavior with fluctuations, but no robust global directionality.

So “direction” in thermodynamics is not something added to time. It is a feature of the distribution of states given special constraints at boundaries. In this view, the arrow is a kind of gradient: a statistical slope from special to typical.

This matters for our taxonomy. It means that “directedness” can arise without a fundamental flow. The arrow is compatible with time-as-coordinate, relational time, or causal order. It is a

property of state space and boundary conditions, not necessarily a property of the parameter used to label change.

7.2 Why records, memories, and prediction align

The thermodynamic arrow is not just about heat and mixing. It is deeply connected to information and to the possibility of stable records. “Records” here means any persistent correlation that can be used to infer something about other events: footprints, photographs, fossils, brain states, notebook entries, ripple patterns, and so on.

Why do records align with the arrow?

Because making a robust record is a physical process that typically involves dissipation and the spreading of correlations into many degrees of freedom. A record is stable when it is redundantly encoded: not held in a single fragile micro-detail, but written into many interacting components such that it can survive noise and perturbation. Redundant encoding is usually achieved by irreversible-looking processes—friction, emission, amplification, scattering—processes that increase entropy in the environment even as they create structured correlations.

This yields a triad of alignment:

- **Memories align with the arrow** because storing memory requires physical substrates that are stabilized by dissipative processes and redundancy.
- **Records align with the arrow** because robust records are created by processes that export entropy, leaving an imprint distributed across many degrees of freedom.
- **Prediction aligns with the arrow** because the direction in which records accumulate is the direction in which agents can condition their beliefs on increasing evidence.

In more direct terms: we know the “past” because the present contains records of it. We do not know the “future” in the same way because the present does not contain stable records of it—except in the weak sense of models and forecasts, which are not records of future events but compressions of regularities.

So the arrow is not only a feature of physics; it is a feature of epistemology embedded in physics. The world’s physical tendency to create stable traces in one direction underwrites the agent’s cognitive asymmetry: remembering versus anticipating.

This also explains why the “experience of flow” is so strongly anchored to entropy. Our conscious sense of “now advancing” is tightly correlated with the fact that our internal state is continually updated by new records arriving and being integrated, while older records remain accessible. The brain is an entropy-exchanging device; its updating is physically situated in the same irreversible macroscopic environment.

7.3 Irreversibility: fundamental or emergent

The word “irreversible” is often used loosely. Physics requires a sharper distinction:

- **Microscopic reversibility (in principle):** if the fundamental laws are reversible and you know the exact microstate, you can reverse the dynamics.
- **Macroscopic irreversibility (in practice):** given only macroscopic information and typical microstates, processes overwhelmingly evolve toward higher-entropy macrostates.
- **Fundamental irreversibility (strong claim):** the fundamental laws themselves are not reversible; information is genuinely destroyed, not merely dispersed into inaccessible degrees of freedom.

Most thermodynamic arrow explanations lean on the second category: emergent irreversibility. It is not that the micro-world forbids reversal; it is that reversal requires extraordinarily precise coordination of micro-details that are effectively never available and are violently unstable under perturbation.

This matters because some metaphysical pictures try to ground “flow” in “irreversibility.” But emergent irreversibility does not automatically grant objective passage. It grants a statistical direction of increasing typicality, a direction in which records become abundant and recoverable.

If one wants irreversibility to be fundamental, one must say what breaks microscopic reversibility and where. That would be a physics claim with potential signatures: deviations from exact reversibility in controlled systems, limits on reconstructing information even in principle, and perhaps new constraints on correlations. But absent such signatures, emergent irreversibility is usually sufficient to account for the arrow we experience.

So the disciplined position is:

- The thermodynamic arrow is real as a macroscopic regularity.
- It can be explained by reversible micro-laws plus low-entropy boundary conditions plus coarse-graining.
- Calling it “fundamental flow” adds extra metaphysical structure unless it predicts new measurable constraints.

7.4 The difference between arrow and parameter

A recurring confusion is to equate “time” with “the arrow.” But these are different categories.

- A **parameter** is a way to label states or events (t, n, a clock variable, a path parameter). Changing the parameterization can leave the physics unchanged.
- An **arrow** is a directionality in state space: a reliable asymmetry that distinguishes one direction along evolution from the other in terms of typicality, entropy, record abundance, and recoverability.

You can have a parameter without an arrow: a system in equilibrium can be parameterized by time while showing no macroscopic directionality. You can also have an arrow defined relative to many different parameters: entropy can increase with respect to a chosen clock variable, or with respect to proper time along a worldline, or with respect to a relational clock, without committing to a single global time coordinate.

This distinction is crucial for our broader project. It means that one can accept a thermodynamic arrow as the primary source of directedness while remaining agnostic about whether the fundamental ordering is time-as-coordinate, relational correlation, causal partial order, or discrete iteration. The arrow is compatible with multiple ontologies.

It also provides a diagnostic against overclaiming: when someone says “time flows because entropy increases,” they may be conflating an emergent arrow (a statistical direction) with a metaphysical flow (a moving present). The arrow explains why one direction is special for records, memory, and action. It does not, by itself, create an objective present that advances.

In the taxonomy of orderings, entropy supplies a powerful form of directionality—an ordering with a slope—without necessarily supplying the parameter that labels progression. That is why it is best treated as an *arrow structure* that can be layered on top of different underlying orderings (correlation-first, causal-first, or even iteration-first), rather than as a full replacement for the ordering parameter itself.

The next section will push this further by examining “records” as an ontology: not just entropy producing an arrow, but present traces being treated as the primary objects from which “past” and “ordering” are inferred.

8. Records Ontology: The World as Consistent Traces

The thermodynamic arrow explains why directionality emerges: low-entropy conditions plus typical evolution produce increasing entropy and abundant records. A records ontology takes a more radical step. It says: perhaps the most fundamental “given” is not an evolving state in

time, but a web of present traces—localized correlations—that stand in mutual consistency relations. What we call the past is then not a thing that exists elsewhere in a timeline, but an inference drawn from the structure of records available now.

This view is often attractive to people who feel that “time flow” is suspect but who also want to take seriously the fact that the world is saturated with evidence: fossils, memories, documents, scars, crater patterns, radiation backgrounds, and so on. A records ontology says: treat those traces as primary. The ordering we experience is reconstructed from how traces cohere.

This section clarifies what a “record” is in physics-friendly terms, how “past” can be inferred from constraints among records, why the persistence of records is nontrivial, and where the records program reaches its explanatory limits.

8.1 Records as localized correlations

A record is a localized physical configuration that is correlated with (and therefore carries information about) other configurations. The key word is correlation. A record is not necessarily a human artifact like a photograph. It can be any stable imprint: a dent in metal, a track in a chamber, a pattern in sediment, a state of a neuron, a magnetic domain on a disk, or a scattered photon field carrying information about what it interacted with.

Two features matter.

- **Locality:** records are embodied in specific subsystems. They occupy regions, substrates, or degrees of freedom that can be pointed to and interacted with.
- **Correlation content:** a record is informative only insofar as it is linked by reliable correlations to other events or states. A random blob is not a record; a structured blob that covaries with something else is.

In an ordinary “time-evolution” picture, records are produced by interactions that spread information from one subsystem into another. In a records ontology, we can flip the interpretive direction: the world “is” a set of correlated subsystems, and the correlations themselves are the primary facts. “History” is then the pattern you infer when you look for a coherent causal/temporal explanation of those correlations.

This move turns the emphasis away from a privileged global parameter and toward the architecture of information in the present. In practice, that architecture is what agents actually have access to. We do not access “the past” directly; we access present records and interpret them.

8.2 “Past” as an inference from present constraints

How can “past” be an inference rather than a thing? The core idea is constraint satisfaction.

Suppose you are given a collection of records: memory states in brains, documents, geological strata, electromagnetic signals, and so on. Many of these records refer to each other: a newspaper matches a dated photograph; the photograph matches a person’s memory; the memory matches the physical scar; the scar matches the medical record; the medical record matches the billing record; and so forth. The records form a mutually reinforcing network.

From this network, you can infer a consistent story: an arrangement of events that would generate the observed pattern of correlations. The “past,” in this sense, is a model that compresses and organizes the record network.

This resembles what science does routinely. We infer dinosaurs from fossils, stellar processes from spectra, early cosmology from background radiation, and personal biography from documents and memories. A records ontology universalizes that stance: all temporal talk is, at bottom, talk about what record networks constrain.

There are two important subtleties:

- **Underdetermination:** many different “histories” might fit a finite set of records. The more redundancy and cross-linking among records, the more constrained the reconstruction becomes.
- **Internal consistency as a criterion of “reality”:** in a records-first view, what makes a past “real” is not that it exists somewhere else in time, but that it is the best-fitting, most coherent, most compressive explanation of the records.

This is an epistemic move with ontological ambition. It says: perhaps the world’s temporal structure is nothing over and above the record constraints that make certain histories the natural compressions of the present.

In the taxonomy, records ontology is a way of generating “ordering” without positing either time-as-becoming or global iteration. It offers an alternative primitive: present informational structure.

8.3 The stability problem: why records persist

A records ontology cannot stop at “records exist.” It must explain why records are stable enough to function as constraints. This is the stability problem.

Records are physical. Physical systems are subject to noise, thermal agitation, perturbations, and decay. Why don’t records instantly dissolve? Why do correlations persist long enough to support memory, science, and the felt solidity of history?

The answer, in physics-friendly terms, is that robust records are typically:

- **Redundant:** the same information is copied into many degrees of freedom. Destroying one copy does not erase the record.
- **Energetically protected:** stable configurations can correspond to minima or barriers in an energy landscape, making them resistant to small disturbances.
- **Environmentally “backed up”:** interactions with the environment spread correlation outward (scattering, emission, diffusion), making records hard to erase globally because erasure would require coordinating an enormous number of degrees of freedom.
- **Coarse-grained stable:** while micro-details fluctuate, macro-patterns persist. A footprint’s microstructure changes, but the footprint remains recognizable.

This is where thermodynamics and records meet. The same processes that drive entropy increase also tend to generate stable records: dissipative amplification and spreading of correlations produce redundancy. The environment becomes a repository of traces.

But there is a deeper point: stable records are not generic. Generic states are high entropy and do not carry interpretable, compressive information about specific structured events. The fact that the world is rich in interpretable records is itself a sign of special conditions—ultimately tied to low-entropy structure and the existence of durable, metastable matter configurations.

So the stability problem pushes records ontology back toward the thermodynamic arrow: records persist because the universe has regimes where metastability, redundancy, and dissipation produce durable correlations. Records ontology is therefore not a free-standing replacement for the arrow; it is often a re-centering of the arrow’s informational consequences.

8.4 What records can explain and what they cannot

Records ontology has real explanatory power, but it also has limits. Naming those limits is essential to keeping the framework disciplined.

What records can explain well:

- **Why the past is knowable:** because the present contains traces that constrain reconstructions.
- **Why memory is asymmetric:** because stable records accumulate in one direction under ordinary macroscopic conditions.
- **Why “history” feels objective:** because independent record streams cross-validate each other, making certain reconstructions overwhelmingly constrained.
- **Why the present feels privileged to agents:** because agency operates on available records and updates; “now” is where records are accessed and integrated.

What records cannot explain by themselves:

- **Why there is any regular lawlike structure linking records:** a network of correlations invites a generative story (dynamics, causal mechanisms). Records alone do not automatically provide the rule that connects them; they only constrain which rules could fit.
- **Why records are produced rather than merely given:** unless one adds a mechanism, a records ontology risks being static: it describes what is, but not how it comes to be. It can say “the world is a consistent record web,” but it must still explain why those consistencies have the particular form we call physical law.
- **Counterfactual dynamics:** physics is not only about reconstructing what happened; it is about predicting what would happen under altered conditions. A pure records stance must show how counterfactuals are grounded in record structure, not merely narrated.
- **The felt flow as a physical primitive:** records can explain the asymmetry of memory and evidence, but whether that fully accounts for the phenomenology of passage remains contested. Records may explain why flow is experienced, without implying that flow is a fundamental physical ingredient.
- **Uniqueness of the past:** finite records can underdetermine history. In practice, redundancy can make histories extremely constrained, but the philosophical point remains: “the past” becomes a best-fit model rather than a uniquely existing domain, unless additional principles are introduced.

In the overall paper, records ontology plays a clarifying role. It shows how much temporal talk can be reconstructed from present informational structure, without invoking a moving present or a universal tick. It also reveals the boundary between explanation and reinterpretation.

Records can reframe the meaning of “past,” but they do not automatically supply the dynamical and counterfactual machinery that physics typically demands.

That sets up the next move: gauge time and reparameterization symmetry. If records and correlations carry the physical content, then the choice of “time parameter” may be closer to a gauge choice than a fundamental fact. The next section will sharpen that idea and show how it collapses many time/iteration disputes into questions about invariance and clock fixing.

9. Gauge Time and Reparameterization Symmetry

A recurring pattern in foundational debates is that people argue fiercely about what “time really is” while quietly assuming that the symbol they use to label change is itself an observable. But in many serious physical descriptions, the parameter that labels evolution is not directly measurable. It is a coordinate choice, and different choices can describe the same physical content. When that happens, “time” behaves like a gauge: a representational redundancy. The physics lives not in the parameter, but in invariant relations among observables.

This is the cleanest way to make sense of why so many time/iteration debates fail to terminate: rival stories often differ only in how they parametrize the same relational structure. This section lays out the idea of time-as-gauge, explains how a physical clock choice “fixes” the redundancy, clarifies what it means for theories to be empirically equivalent, and shows how “iteration” frequently collapses into a gauge choice once its hidden assumptions are removed.

9.1 Time as gauge: why parameter choice can be unphysical

Gauge, in the broad sense used here, means a freedom to change the mathematical description without changing any observable prediction. When a theory has gauge freedom, multiple different-looking descriptions correspond to the same physical situation. The classic lesson is: do not mistake the gauge-dependent elements of the description for physical entities.

Time-as-gauge appears when the equations that define the physical content are invariant under reparameterizations of the evolution label. Informally: if you can replace the parameter t with any smooth, monotone function $t \rightarrow t'(t)$ and rewrite the equations in the new parameter without changing the predicted relations among observables, then the parameter itself is not physical. It is a label.

In such a setting:

- The parameter is a bookkeeping device.
- Physical statements must be phrased in parameter-invariant form.

- “What time is” becomes the question: which relational structures remain invariant when you change the parameter?

This does not mean “nothing happens.” It means that “happening” is not tied to a unique external clock variable. The theory describes allowed correlations among variables. A parameter is introduced to solve, simulate, or visualize those correlations, but different parameterizations are equally valid.

The key discipline is: if changing the parameter changes nothing observable, the parameter is not an element of reality in the same sense as an observable. It is a coordinate on the space of descriptions.

This viewpoint aligns naturally with correlation-first ordering. If all we ever measure are correlations between physical subsystems, then any extra label we attach that can be reshuffled without consequence is surplus structure.

9.2 Fixing time by choosing a clock: what becomes real

If time is gauge-like, how do we ever do concrete physics? We do it by *fixing the gauge*—by choosing a clock.

Choosing a clock means selecting a particular physical degree of freedom C and using its reading to parameterize the rest: you define “time” operationally as “what that clock indicates.” This converts the abstract, redundant parameter into an empirically anchored variable.

Two things happen when you do this:

- **You obtain a usable evolution description:** you can write $S(C)$ for the system variables S as functions of the clock variable C . Predictions become statements like “when the clock reads $C = c_1$, the system will (with such-and-such probability) have S in this range.”
- **You declare what counts as the time coordinate for your experiment:** not as a metaphysical absolute, but as a convention grounded in a stable subsystem.

What becomes real is not the bare parameter, but the *correlation structure relative to the chosen clock*. The physical content is: given the clock’s state, what do other observables do? The reason this feels like ordinary time is that good clocks are stable, portable, and widely shareable; their correlations with many systems are simple and reliable.

This also clarifies why multiple clocks are not a conceptual disaster. Different clock choices are different gauge fixings. If the framework is coherent, changing clocks should change the intermediate description but not the invariant predictions.

So “time” becomes real in a qualified way: not as a unique cosmic parameter, but as an operational coordinate derived from a chosen physical subsystem, with reality residing in the invariants that survive clock choice.

9.3 Empirical equivalence: when rival stories are the same theory

Two frameworks can feel philosophically different while being empirically equivalent: they predict the same outcomes for all possible measurements (within their shared domain). In such a case, physics cannot choose between them by experiment; the choice becomes one of interpretation, convenience, or metaphysical taste.

Time debates frequently live here. Consider three descriptions:

- A time-as-coordinate description with t as an external parameter.
- A relational description with an internal clock C and other variables expressed as functions of C .
- An iterative description indexing states by n .

If all three are connected by reparameterizations and clock selections that preserve the predicted correlations among observables, then they are not competing physics theories. They are different representational skins on the same predictive core.

The practical significance is not trivial, though. Even if empirically equivalent, different stories can:

- Highlight different invariants.
- Suggest different generalizations.
- Make different problems appear natural (for example, focusing on constraints versus evolution).
- Encourage different intuitions about what would count as “new physics.”

So empirical equivalence does not mean “no difference.” It means “no measurable difference given current definitions of measurement.” For physics, that puts a ceiling on claims. You may prefer one ontology, but you cannot claim it is *determined* by experiment unless it yields inequivalent predictions.

This is exactly the discipline we need for the paper: separate what can be empirically decided from what is interpretive overlay.

9.4 How “iteration” often reduces to gauge choice

Once you adopt time-as-gauge and clock fixing, a large fraction of “iteration” talk collapses into one of two cases.

Case A: Iteration is reparameterization.

You label the evolution by an index n instead of a parameter t . Provided n can be mapped monotonically to t (or to a clock variable C), and all observable correlations are preserved, iteration is a gauge choice. It does not add discreteness; it adds a different label.

Case B: Iteration is a clock choice.

Sometimes “iteration” really means “counting cycles/ticks of a particular clock.” In that case, n is not a metaphysical update counter. It is an operational variable: the number of oscillations, the number of transitions, the number of counted events. That is simply relational time in a discrete readout form.

In both cases, iteration is not a rival ontology to time; it is either a different parameterization or a specific clock-based gauge fixing.

What remains as genuinely distinct is only the stronger iteration claim: a *true tick* that cannot be transformed away, or a computational update rule with resource bounds that yields new constraints. Those are not gauge choices; they are new structure. But note what must happen for them to be truly distinct: they must break or restrict the reparameterization freedom in a way that shows up as invariant, measurable effects.

This gives a clean sorting rule:

- If “iteration” can be removed by reparameterization, it is gauge.
- If “iteration” is just counting a clock’s ticks, it is clock choice.
- If “iteration” asserts a fundamental discrete update that changes invariants or predictions, it is a physical claim and must be tested or constrained.
- If “iteration” implies global synchronization, it likely smuggles in a preferred structure and must explain why that structure is not already visible.

So the gauge-time viewpoint is not an attempt to dissolve debate by fiat. It is a filter that tells us which debates are about physics and which are about representation. It preserves the possibility that iteration is fundamental, but it refuses to let “iteration” function as a metaphysical badge unless it comes with consequences that survive changes of parameter and clock.

This sets up the next stage of the paper: once we recognize that time-parameters can be gauge-like, we can explore other orderings—phase, scale, inference, agency—that also function as ways of organizing correlations without insisting on a single universal flowing time.

10. Phase Ordering: Cycles, Angles, and Return

If time-as-coordinate suggests a line, and iteration suggests a sequence of steps, phase ordering suggests a circle. Many physical systems are not best understood by an ever-advancing parameter but by a repeating cycle: oscillators, rotations, waves, orbital motion, heartbeat rhythms, circadian patterns, and countless engineered feedback loops. In such systems, what matters is often not “how much time has passed” in an absolute sense, but “where you are in the cycle.” The relevant ordering variable is an angle-like quantity: phase.

Phase ordering becomes philosophically interesting because it offers a way to define “before/after” locally without demanding global linear progression. It also clarifies why “clocks” so often are oscillators: a clock is frequently a physical system whose state winds around a cycle in a stable way. But phase is not the same as time. Phase repeats. It supports ordering within a cycle, but it does not uniquely encode long-run progression unless additional structure is added (cycle count, drift, or an auxiliary monotone variable).

This section explains phase as an ordering principle, distinguishes physical cycles from cosmological cycles, shows how phase can replace linear time in certain contexts, and carefully separates phase-based ordering from record-based ordering.

10.1 Periodic variables as order parameters

An order parameter, in the broad sense used here, is a variable that captures the salient structure of a system’s behavior and can be used to organize or predict other variables. For periodic systems, phase is such a variable.

Consider a stable oscillator. Its detailed state may be high-dimensional (positions, momenta, internal degrees of freedom), but much of its behavior can be compressed into a phase variable ϕ that advances around a circle. When ϕ increases smoothly, you can predict where the system is in its periodic cycle: crest vs trough, systole vs diastole, rising edge vs falling edge.

Phase provides ordering in a very specific sense:

- Within a cycle, phase gives a “before/after” relation: $\phi = 0.2$ precedes $\phi = 0.3$ along the direction of phase advance.
- The ordering is modular: $\phi = 0.9$ is “near” $\phi = 0.1$ because the cycle wraps.

- Phase ordering can be highly robust: even when amplitude fluctuates or noise perturbs the system, the phase can remain coherent and predictive.

This is why phase is a natural “internal time” variable in many contexts: it is a coordinate on a recurring process that can parameterize other changes.

But phase alone is not a complete substitute for linear time. Because phase wraps, it cannot distinguish “one cycle later” from “this cycle now” unless you also track a winding number (cycle count) or use a slowly varying auxiliary variable. That fact will matter when we ask whether phase ordering could be “fundamental” rather than merely useful.

10.2 Cyclic cosmologies and effective cyclic dynamics

Phase ordering tempts people toward cyclic cosmologies: perhaps the universe itself runs on cycles, returning through epochs. The conceptual attraction is obvious: it replaces one-way progression with recurrence and can make “beginning” and “end” feel less absolute.

But the paper’s discipline requires separating three very different claims:

- **Local cyclic dynamics:** specific subsystems oscillate (this is uncontroversial and ubiquitous).
- **Effective cyclic behavior:** a complicated system exhibits approximate cycles in some reduced variables even if the full microstate does not repeat.
- **Cosmic cyclicity:** the entire cosmos undergoes large-scale recurrent phases.

Only the first is directly supported as a general principle by everyday physics. The second is common in complex systems: macro-variables can look cyclic while micro-details never return. The third is a strong claim that would require a particular global structure and constraints on entropy, record accumulation, and the persistence of information across cycles.

The key conceptual tension is: cycles imply return, but records and entropy tend to accumulate. If the universe truly repeats in a strong sense, then one must explain either:

- How entropy is reset or exported in each cycle, or
- How the cycles are not exact repeats but only repeats in coarse-grained variables, or
- How “return” should be understood when microstates differ.

This tension is not an argument against cyclicity; it is a sorting mechanism. Many “cyclic” narratives are really about effective cycles in some macro variables, not literal recurrence of the total microstate.

So, phase ordering is compatible with cycles at multiple levels. The question is always: what exactly repeats, at what fidelity, and how does that coexist with the arrow-like behavior associated with stable records and entropy gradients?

10.3 When a phase variable replaces linear time

Phase can replace linear time as the primary ordering variable when the phenomena of interest are dominated by periodic structure and when measurements naturally reference cycle position rather than elapsed duration.

This happens in at least three broad situations:

- **Synchronization problems:** When you care about coordination between oscillators (two pendulums, coupled lasers, neural rhythms), the key variable is phase difference, not absolute time. “Who leads” and “who lags” are phase relations.
- **Control and feedback:** In many engineered and biological systems, control is phase-locked: actions depend on cycle position (timing of heart pacing, gait, alternating current power, phase-locked loops).
- **Reduction of dynamics:** When a system has a stable limit cycle, many perturbations decay back to the cycle. The long-term behavior can be parameterized mainly by phase along the cycle.

In such contexts, phase ordering is not merely a convenient representation; it is the variable that captures what is robust and predictive. Linear time can still exist as a background coordinate, but it is secondary. The system’s state is better tracked by where it is in its cycle.

Now, could phase ordering ever be fundamental, not just local? Possibly, but it would require additional structure, because phase is bounded and repeats. A fundamental ordering must support the accumulation of novelty, records, and non-repeating structure—or else reinterpret those phenomena as emergent patterns on top of recurring base dynamics.

One way this can happen is through two-layer structure:

- A fast cyclic variable (phase) supplies local ordering and coordination.
- A slow drift variable supplies long-run progression, novelty, or accumulation.

In that composite picture, the world is “phase-ordered” at one level and “arrow-ordered” at another. Many real systems already behave this way: fast oscillations embedded in slow development (daily cycles embedded in aging, business cycles embedded in secular growth, neural oscillations embedded in learning).

This ties back to our initial question: “other than time and iteration, what orderings exist?” Phase is one such ordering, but it often wants a companion variable to account for irreversibility and record accumulation.

10.4 Distinguishing phase-order from record-order

Phase order and record order can be confused because both allow you to infer “where you are” from local data. But they are fundamentally different in what they treat as primary.

Phase-order is about position in a cycle. It is intrinsically periodic. It supports ordering within recurrence.

Record-order is about traces that encode constraints about other events. It is intrinsically cumulative: records tend to persist and cross-link, producing a directionality tied to stability and redundancy.

Here are the crisp differences:

- **Topology:** phase lives on a circle; record structure typically lives on a directed, branching network of constraints (often aligned with an arrow).
- **Repetition:** phase repeats by design; records accumulate and, in practice, do not “wrap around” without deliberate erasure.
- **Predictive content:** phase tells you what comes next within the cycle; records tell you what likely happened to produce the current constraints and what constraints will likely persist.
- **Reset:** phase resets every cycle; records do not reset automatically and often require energetic work to erase or overwrite.
- **Directionality source:** phase direction is often defined by convention (increasing angle); record direction is defined by the physical asymmetry of record formation and stability.

In many systems both are present: a clock is phase-based (oscillation) and also record-based (counted ticks recorded in memory or a counter). This hybrid is what makes clocks powerful: they combine repeatable cycles with cumulative counts.

So, phase ordering is a distinct element in the taxonomy: it provides an ordering principle rooted in periodicity and synchronization rather than in linear progression. But it does not, by itself, explain the one-way accumulation of evidence and memory that dominates our lived experience. That role is more naturally played by entropy and records. In the larger map, phase ordering often sits as a local, cyclic coordinate system layered inside a broader arrow-structured world.

The next section will shift from cyclic ordering to scale ordering: a different kind of “before/after” defined not by cycles or by time, but by moving between levels of description from microscopic to macroscopic via renormalization-like flow.

11. Scale Ordering: Renormalization as a Kind of “Before/After”

So far we have treated ordering as something that organizes events (time coordinates), influences (causal order), evidence (records), directionality (entropy), repetition (phase), or representation freedom (gauge time). Scale ordering introduces a different axis entirely: an ordering not of *moments* but of *descriptions*. The guiding idea is that physics is layered. What you call “the same system” can be described at many resolutions, from microscopic details to macroscopic variables. Moving between these layers is not merely a matter of convenience; it often reveals structural stability—universal behavior that does not depend on micro-details.

Renormalization is the disciplined machinery that formalizes this movement between scales. For our purposes, its philosophical significance is that it provides a kind of “before/after” that is neither time-like nor iterative in the usual sense. It is a flow on the space of descriptions: “from fine to coarse,” “from micro to macro,” “from detailed to effective.” That flow can be treated as an ordering principle because it is directional (coarse-graining discards detail), compositional (you can coarse-grain in stages), and predictive (it tells you which macro-behaviors are stable under variation of micro-structure).

This section explains scale ordering, why it behaves like a flow, what it illuminates about universality and stability, and why it cannot, by itself, replace time as an ordering of events.

11.1 Ordering by description scale, not moments

To order by scale is to order by *resolution*. Instead of asking “what happens next in time,” you ask “what remains when I stop caring about details below some cutoff.” The operation is: take a complex system with many degrees of freedom and form a reduced description that tracks only the variables relevant at a chosen scale.

Two key asymmetries make this an ordering rather than a reversible mapping:

- **Information loss:** coarse-graining discards micro-details. Many microstates map to the same macrostate. This makes the transformation many-to-one.
- **Irreversibility in description:** because details are discarded, you cannot uniquely recover the fine description from the coarse one without adding extra assumptions.

This makes scale ordering directional in a way that resembles an arrow: fine → coarse is natural and compositional; coarse → fine is underdetermined.

It is important to see what is being ordered. It is not the world “moving” in time. It is our description moving along a hierarchy of models. The ontology question becomes: is this merely epistemic (about what we know and ignore) or does it reflect something deeper about how physical law organizes itself across scales?

Renormalization suggests the second: scale transformations are not arbitrary ignorance; they expose structural features of physical systems. Many large-scale behaviors are not sensitive to micro-details because they are controlled by stable structures in the space of theories. That is a profound fact about the world, not just about our limitations.

11.2 Effective theories and emergence as a flow

An effective theory is a model that is accurate within a given range of scales and energies while ignoring fine details that do not matter at that range. The renormalization viewpoint treats “emergence” as a flow: as you move to larger scales, the parameters that describe interactions “run,” and different micro-theories can converge to the same effective behavior.

This yields a “before/after” metaphor that is surprisingly disciplined:

- “Before” = the micro-description with many degrees of freedom and specific interaction details.
- “After” = the macro-description where irrelevant details have washed out, and only certain combinations of parameters matter.

The term “flow” here is not time flow. It is flow in a space of descriptions. But it behaves like a dynamical process: repeated coarse-graining steps can be composed, and the resulting effective description can approach stable forms.

This gives emergence a formal backbone. It is not merely “the whole is more than the parts.” It is: there are stable macro-regularities that persist under wide variation in micro-details. That stability can be described as an attraction phenomenon in theory-space.

This is why scale ordering is relevant to our paper. It offers another way to talk about “order” that does not depend on time or iteration. You can ask: what is stable as we change scale? What patterns survive? What variables become the right coordinates? In this sense, “order” is what the world looks like when you move upward through layers of effective description.

11.3 What scale-order explains about stability and universality

Scale ordering is especially powerful at explaining two things that time-centered metaphysics often takes for granted: stability and universality.

- **Stability:** Many macroscopic phenomena are stable under perturbations. Small micro-changes do not radically alter macro-behavior. Scale ordering explains this by showing that many micro-details are “irrelevant” to the macro-description: they do not survive coarse-graining as controlling variables.
- **Universality:** Very different microscopic systems can share the same macroscopic behavior. This is the famous surprise: the details differ, but the large-scale behavior falls into a small number of classes. Scale ordering explains universality by identifying why multiple different micro-models flow toward the same effective description.

In the language of our taxonomy, scale order is an ordering of “what matters.” At fine scales, many variables matter. At coarse scales, fewer do, and they are often collective variables. This is an ordering principle that compresses reality.

Scale ordering also clarifies the status of many “fundamental” debates. Often, what we experience as deep is an emergent macro-regularity that is insensitive to micro-details. If so, insisting on a micro-ontology to explain the macro can be misguided. The right approach is to identify the stable macro invariants and show how they arise under scale flow.

This has a direct connection to time and iteration:

- If you propose a discrete iterative micro-world, scale ordering asks: does your micro-structure flow to the observed smooth macro-symmetries? If yes, discreteness can be hidden. If no, discreteness is constrained.
- If you propose time-as-becoming, scale ordering asks: is becoming a micro-primitive, or is it an emergent agent-level phenomenon that appears at certain scales (where memory, control, and records exist)?

So scale ordering becomes a diagnostic tool: it tells you where to locate claims—micro, meso, macro—and whether they should be expected to be robust or fragile.

11.4 The limits: why scale-order is not a full substitute for time

Despite its power, scale ordering cannot replace time as the ordering of events. The reason is simple: scale ordering orders descriptions, not occurrences.

Time (or causal order) answers questions like:

- Which events can influence which?
- What happens next if I act now?
- How do signals propagate?

- How do records accumulate?

Scale ordering answers different questions:

- Which variables matter at this resolution?
- Why are some macro-laws universal?
- Why do different micro-systems look the same at large scales?
- What patterns survive coarse-graining?

You can have scale flow in a static world, because it is a map between descriptions, not an account of temporal change. Even if you treat the renormalization flow as “dynamics in theory-space,” it does not tell you how a particular system evolves from one state to another. It tells you how your model changes when you change what you count as relevant detail.

This distinction prevents a category error: using scale as a poetic replacement for time. Scale ordering is real and deep, but it is orthogonal to temporal ordering. At best, it can complement time by explaining why time-based laws take the forms they do at different scales, and why certain temporal behaviors (diffusion, dissipation, wave propagation) are robust across micro-implementations.

In the larger architecture of the paper, scale ordering provides a second axis of “ordering” that is not about before/after in the lived sense, but about fine/coarse in the descriptive sense. It helps explain why the world is intelligible at all: because macro-regularities are stable and compressible. But it does not by itself generate the directed sequence of experience, memory, and action. For that, one needs either a time/clock parameterization, causal partial order, entropy/records, or an agency-centered update sequence.

The next sections will push into inference/computational depth and agency ordering—two more “orthogonal orderings” that, like scale ordering, organize reality without necessarily being the time parameter itself.

12. Inference and Computational Depth as Ordering

Time orders events. Causal order orders influence. Entropy and records order evidence and directionality. Scale orders descriptions by resolution. Inference and computational depth propose something even more abstract: perhaps “before/after” can be grounded in the structure of *derivation*—what can be computed from what, how hard it is to derive, how compressible a description is, and how much “computational work” separates one state from another.

This family of ideas is attractive because it connects ordering to intelligibility. Our strongest sense of direction often attaches to notions like “building,” “forming,” “learning,” “organizing,” and “elaborating.” Those can be framed as increases in computational depth: the system’s present state encodes a history of constraints, computations, or selections. In this view, “later” is the direction toward greater structured complexity or greater depth of derivation.

But this move is dangerous if handled loosely. Complexity can mean many things, and derivability can be purely epistemic: it can reflect what we can compute, not what the world does. The section therefore does two jobs: it lays out how inference-based orderings can be made precise enough to be meaningful, and it marks the boundary where they collapse into a story about knowledge rather than a claim about reality.

12.1 Ordering by derivability, proof depth, or compressibility

The most basic inference-order idea is: define an ordering relation between states by whether one can be derived from another under some rule set.

- If state B can be derived from state A by applying lawful transformations, then A precedes B in derivation order.
- If deriving B from A requires a long chain of transformations, B has greater proof depth relative to A.
- If B can be summarized by a shorter description than A (or vice versa), you might define an order by compressibility: the “simpler” description is more basic.

There are multiple ways this can be instantiated:

- **Derivability order:** $A \rightarrow B$ if B follows from A under the dynamics or constraints.
- **Proof-depth order:** A precedes B if the minimal derivation of B from A has fewer steps (or lower resource cost).
- **Compressibility order:** A precedes B if A is a shorter program/description that generates B (or that generates the distribution containing B).

Notice the immediate ambiguity: which rules count as legitimate derivations? Which cost function measures depth? Which description language defines “short”? These are not trivial details. They determine whether the ordering is objective or merely conventional.

Still, the idea has real traction because physics itself frequently uses derivational structure. We often do not observe microstates directly; we infer them. We often describe systems by compact laws that generate a vast range of behavior. In such contexts, compressibility is not just about our ignorance; it reflects that the world has regularities that admit short descriptions.

So an inference-based ordering is best framed as: ordering by the structure of generative constraints and the minimal resources required to produce or specify what we observe.

12.2 Description length versus physical process

Here is the key tension: a short description does not necessarily imply a short physical process.

A system can be physically “deep” but algorithmically “shallow,” or the reverse, depending on what you mean:

- A simple rule can generate extremely complex-looking outcomes after many iterations. The description is short; the process can be long.
- A complex configuration might be compressible if it has hidden structure or symmetry, even if physically it arose through a long, contingent history.
- Conversely, a configuration might be incompressible (no shorter description than listing it) even if it arose from a simple physical process with randomness.

So we must separate two notions of depth:

- **Descriptive depth:** how long the shortest description/program is.
- **Computational depth:** how much computation is required to generate the object from its shortest description.
- **Physical depth:** how much physical time, energy dissipation, interaction, or causal chaining was required in the actual world to bring it about.

In a purely epistemic frame, these are properties of our representations. In a physics-facing frame, we ask whether there is any principled alignment between them. Sometimes there is:

- Physical laws that are simple can generate physically deep structures over time (e.g., stable patterns formed by repeated local interactions).
- Record formation can encode long causal histories into present states, making some present configurations physically deep: they embody many past interactions.

But the alignment is never automatic. If we declare “the arrow of time is the direction of increasing complexity,” we risk a false generalization because complexity can decrease locally, fluctuate, or be traded between system and environment.

So the disciplined approach is: inference-ordering can be meaningful, but only when it is anchored to physical generative mechanisms or to invariant constraints, not just to our choice of description language.

12.3 When “complexity” functions like an arrow

Despite the cautions, there are contexts where complexity-like measures behave arrow-like: they tend to increase in a robust direction under typical conditions, providing a “before/after” gradient.

There are at least three ways this can happen:

- **Entropy-driven structure formation in open systems:** In systems driven far from equilibrium, local order can increase while global entropy increases. Complexity can grow in the sense of richer patterns and functional organization, even as dissipation continues. This is not a violation of the thermodynamic arrow; it is a complement to it.
- **Accumulation of records:** As records accumulate, the present can encode more constraints about “what had to happen.” The system’s state becomes more “historically laden.” That can be interpreted as increasing computational depth: more causal structure is embedded in the present.
- **Learning and adaptation:** Agents and adaptive systems update internal models. Their state can come to encode more predictive structure about the environment. That is an arrow in model space: a direction of increasing fit, increasing internal organization, increasing usable information.

In these cases, complexity functions like an arrow not because the universe has a primitive “complexity time,” but because typical dynamics plus boundary conditions yield gradients in certain measures of structured organization.

However, the arrow is conditional:

- It depends on openness, driving, and constraints.
- It can stall, reverse, or plateau.
- It is not universal across all subsystems.

So complexity-as-arrow is best treated as a mesoscale phenomenon: an ordering principle that applies in certain regimes (especially where agents, life, or long-lived far-from-equilibrium structures exist) rather than as a universal replacement for time.

This is useful for our taxonomy because it highlights a second kind of directedness besides entropy: an arrow of accumulated structure, accumulated constraints, accumulated learned models. It explains why “later” sometimes feels like “more built” rather than merely “more disordered.”

12.4 Risks: confusing epistemic order with ontic order

The biggest hazard in this section is category error: confusing epistemic order (order in our knowledge or descriptions) with ontic order (order in the world).

Here are the main risks:

- **Language dependence:** compressibility depends on the description language. Without care, “shorter” is not objective.
- **Observer dependence:** what is derivable depends on what rule set the observer uses and what information they possess.
- **Conflating inference with causation:** the fact that we can infer A from B does not mean A caused B. Reverse inference is common: present evidence lets us infer earlier causes, but that does not invert causal direction.
- **Smuggling in time:** many computational notions implicitly assume a step count or a time-like resource. If you define depth by number of steps, you have reintroduced iteration unless you can define step-cost in a parameter-invariant way.
- **Moralizing complexity:** it is tempting to treat “more complex” as “more advanced” or “more real.” Physics does not license that. Complexity can increase, decrease, and redistribute; it is not automatically a universal telos.

To keep the framework physics-facing, the paper will adopt a disciplined stance:

- Use inference/computational depth as an ordering only when it can be anchored to physical mechanisms (record accumulation, causal chain length, resource-bounded propagation, learning dynamics) and not merely to narrative intuition.
- Treat complexity gradients as regime-dependent, not universal.
- Keep clear separation between “the world’s causal structure” and “our inferential reconstruction of it.”

With these safeguards, inference-based ordering becomes a valuable member of the taxonomy. It explains why “order” can mean “derivable from,” “compressed by,” or “built through” in contexts where time/iteration language feels too crude. But it also marks a boundary: without operational anchors, inference ordering is about how we know, not about what is.

The next section will make that boundary even more concrete by turning to agency and decision ordering, where the ordering variable is explicitly tied to information availability and control-loop updates—structures that live at the interface between epistemology and physics.

13. Agency and Decision Ordering

Agency-ordering treats “before/after” as something defined inside the loop of acting systems. Instead of starting with an external parameter (t) or a global tick (n), it starts with a structural fact about agents: they take in information, update internal state, choose actions, and thereby change the stream of future information. The ordering variable is not “time” in the abstract; it is *availability*: what information is accessible to the agent at the moment of choice, and what actions are feasible given that information.

This framing is not a retreat into psychology. It is a way to describe a physically embedded process: control systems exist as subsystems of the world, constrained by signaling, locality, noise, memory, energy, and the thermodynamic cost of maintaining organized state. Agency-ordering therefore sits at an interface: partly epistemic (beliefs, updates), partly ontic (causal constraints, record formation, resource limits). It offers a disciplined explanation for why experience feels sequenced and directional even when “time” as a metaphysical primitive is left undecided.

13.1 Control loops: order defined by information availability

A control loop is the minimal skeleton of agency:

1. sense (receive measurements),
2. estimate (compress measurements into a state),
3. decide (select an action),
4. act (apply an intervention),
5. observe the consequences, and repeat.

The crucial ordering relation here is *what was known when the decision was made*. That yields an intrinsic “before/after” even without a global time parameter:

- **Before** = the agent has not yet received a particular measurement, or has not yet integrated it into its internal state.
- **After** = the agent has received and integrated the measurement, and can condition decisions on it.

This is an ordering by information sets. The agent’s internal state partitions the world into “distinctions the agent can currently act on” and “distinctions not yet available.” Decisions are ordered by which distinctions were available at the decision point.

Two consequences follow immediately.

First, control loops create a natural partial order: two agents (or two subsystems) can be at different “information stages” with no need for global simultaneity. They can each have their own notion of “now” defined by their most recent sensor integration and action selection.

Second, control loops formalize the sense in which “the present” is operational: it is the slice of the world that has been incorporated into the agent’s actionable state. The “present” is not a cosmic hyperplane; it is a boundary between unintegrated signals and integrated state.

Agency-ordering also makes the role of locality explicit. Information arrives at finite speed. Sensors have spatial placement. Signals have delays. Thus, what an agent can know “now” is bounded by causal horizons. In this view, agency is causality made practical: an agent’s loop is a local causal engine that turns accessible correlations into interventions.

13.2 Belief update sequences and the arrow of learning

Belief updates define a second agency-based ordering: the sequence of model revisions as evidence accumulates. Even if microphysics is reversible in principle, learning is effectively arrowed because:

- evidence arrives in one direction (you acquire new records),
- assimilation is physically implemented in memory substrates (which are stabilized by dissipative processes),
- forgetting and unlearning are not symmetric inverses of learning (they require specific interventions and often degrade detail).

So, for an agent, “later” is naturally “after more updates.” The ordering variable is the update index k in a sequence of internal states:

$q_0 \rightarrow q_1 \rightarrow q_2 \rightarrow \dots$

where each transition corresponds to a belief update triggered by new data.

This is not automatically the same as a universal iteration index for the cosmos. It is local, agent-relative, and physically realized. But it produces a robust experiential arrow:

- **Prediction improves (or at least changes) in a directed way** because the agent conditions on an expanding set of records.
- **Narrative coherence increases** because the agent’s internal state increasingly compresses the record stream into stable causal expectations.

- **Commitment becomes path-dependent** because once a model has been updated, the next update is conditioned on the current model, not the original one.

Agency-ordering therefore gives a precise interpretation to “the arrow of learning”: it is the direction along which the agent’s internal compressions and control policies are revised under incoming evidence. It is also inherently resource-bounded: agents cannot store or compute everything, so they form compressed models. That compression creates hysteresis and irreversibility at the level of belief, even in a world where micro-dynamics may be reversible.

This is where agency-ordering connects back to earlier sections:

- It aligns with **records ontology**: learning is the internalization of records into controllable state.
- It aligns with the **thermodynamic arrow**: stable learning requires dissipation and robust memory substrates.
- It aligns with **inference/computational depth**: learned models can be seen as compressed descriptions that encode deep regularities.

So the arrow of learning is not a competitor to the thermodynamic arrow; it is one of its most visible local manifestations inside an information-processing subsystem.

13.3 Multi-agent consistency: when orders conflict

Once you admit multiple agents, you immediately face a problem that looks like “time disagreement” but is actually “order disagreement”: different agents have different information sets, different clocks, different update schedules, and different action horizons. Their internal orders can conflict.

There are three primary conflict modes.

- **Asynchronous evidence arrival**
Two agents receive relevant data in different sequences due to locality and signaling delay. They update in different orders, leading to different intermediate beliefs even if they eventually converge.
- **Different model classes and priors**
Even with identical evidence, agents may compress differently. One agent treats a pattern as signal; another treats it as noise. Their belief-update orderings diverge not because the world differs, but because their internal representational constraints differ.
- **Strategic interaction**
In multi-agent environments, agents do not merely observe the same world; they

observe each other. This creates reflexive loops: my action changes your observations, which changes your next action, which changes my observations. Order becomes entangled with game dynamics.

In this setting, “consistency” is not automatic. It must be engineered or earned through protocols: shared measurement standards, synchronization conventions, communication, cross-validation, and conflict resolution mechanisms. This is the multi-agent analogue of the records cross-linking principle: robust agreement arises when independent channels constrain the same latent structure.

The key conceptual payoff is that agency-ordering makes disagreement about “what is happening now” unsurprising. It is not primarily metaphysical. It is structural: different agents occupy different causal positions and have different information boundaries.

This also offers a disciplined way to discuss “presentness” without global simultaneity. Each agent has a present defined by its actionable information set, and the relationship between presents is mediated by causal signaling and record exchange.

13.4 What agency-order can illuminate about experience

Agency-ordering is especially useful because it can explain major features of lived temporal experience without assuming an objective flowing time.

Four experiential facts fall naturally out of the agency picture:

- **The felt “now” as an update boundary**
“Now” is the interface where new inputs become integrated into the internal state that governs action. It is not a universal slice; it is the agent’s boundary between unassimilated signals and assimilated control state.
- **The felt direction of time as record accumulation inside the agent**
We experience “past” as accessible memory and “future” as open possibility because memory is physically stored and stabilized, while future records do not yet exist in our possession. This is the records/entropy arrow expressed through an agent’s internal architecture.
- **The felt continuity of flow as continuous control**
Agents run control loops at high frequency. Even if the underlying world were discrete at some substrate scale, the agent’s integration and action could be smooth at its operating scale. “Flow” can therefore be an emergent phenomenology of dense update cycles and predictive smoothing.

- **The irreducible first-person asymmetry**

From the outside, an agent can be modeled as a physical system with state transitions. From the inside, the agent's access to its own state is privileged and its access to the world is mediated by sensors. This generates a persistent asymmetry: the world is encountered as a stream of surprises relative to an internal model. That stream-like character is an experiential ordering even if the underlying physics is formulated timelessly.

Agency-ordering therefore serves as a bridge concept in the paper's taxonomy. It shows how "order" can be physically grounded in information availability, local causality, and memory formation, while remaining compatible with multiple choices about fundamental ontology (block time, relational time, causal sets, or discretized dynamics).

It also provides a practical moral: when debates about "time" become stuck, it can be clarifying to ask what structure is actually doing the work in experience and in practice. Often it is not a cosmic time fluid. It is the agent's loop: sensing, updating, acting, recording, and learning under finite resources and causal constraints.

In that sense, agency-ordering does not replace time; it explains why time-talk is so natural and so sticky. It identifies the physical substrate of "sequence" at the scale where sequence matters most to us: the scale of decisions.

14. Empirical and Conceptual Discriminators

Up to this point, the paper has widened the menu of "orderings of change" without insisting that physics must pick a single winner. This section turns the crank in the opposite direction. It asks: if two stories really differ, what would force a difference in what can be observed, inferred, or operationally constrained? And if two stories are empirically indistinguishable, how should we talk about them without inflating metaphysics into faux measurement?

The goal is a set of discriminators that separate four cases: (i) pure re-labelings, (ii) different clock choices, (iii) genuinely different physical structures that could, in principle, be tested, and (iv) metaphysical add-ons that may be meaningful but are not physics unless they couple to observables.

14.1 What would count as evidence for fundamental discreteness

Fundamental discreteness is not "quantization of measurement outcomes" or "digital simulations approximate well." It is the claim that the world's underlying evolution or spacetime structure has a smallest meaningful step, grain, or event atom that cannot be refined away.

Evidence for discreteness would look like an invariant residue that survives all attempts to reinterpret it as noise, finite instrumentation, or a coordinate choice. In practice, that means one or more of the following patterns:

- A resolution limit that is not merely technological: as instruments improve, the same floor remains in a way that can be modeled as intrinsic rather than environmental.
- Systematic deviations from smooth continuity assumptions that depend on scale in a predictable way (not random drift, not calibration error).
- A universal “jitter” or granularity signature that appears across multiple experimental platforms that share no common hardware weakness.
- Discreteness that implies a maximum meaningful frequency or minimal time interval in an operational sense, not just a philosophical sense.

A crucial caveat is that discreteness does not automatically force obvious artifacts. Discrete substrates can produce extremely smooth emergent behavior, just as molecular matter produces smooth fluids. So the discriminator is not “do we see steps?” It is “does the theory require a non-removable scale that leaks into observable correlations in at least some regime?”

14.2 Lorentz invariance, dispersion, anisotropy, and noise floors

Many discreteness pictures accidentally smuggle in a preferred structure: a global tick, a lattice direction, or a special slicing of events. The most sensitive discriminators are therefore those that look for tiny breaks in symmetry or tiny distortions in propagation.

Four families of signatures recur across otherwise different models:

- Lorentz invariance stress tests
If there is a preferred frame or preferred foliation, then the laws can acquire small frame-dependent corrections. Operationally this would show up as direction- or velocity-dependent deviations in otherwise invariant relationships. The discriminator is not philosophical preference; it is whether any measurement outcome depends (even slightly) on absolute motion relative to some hidden structure.
- Dispersion stress tests
Smooth relativistic propagation often implies a specific relationship between frequency (or energy) and speed of signals. Many discrete or granular models predict a modified relationship, typically making propagation slightly energy-dependent or frequency-dependent. Over long baselines, tiny differences can accumulate. A true discriminator would be a reproducible, energy-correlated delay or phase drift that cannot be attributed to known media effects or source mechanisms.

- Anisotropy stress tests

A lattice-like substrate, a preferred axis, or an emergent grid can produce tiny directional asymmetries: the same experiment rotated in space yields slightly different outcomes. The key here is repeatable orientation dependence that tracks sidereal or inertial orientation rather than local environment.

- Noise-floor stress tests

If evolution proceeds in discrete updates, there may be an irreducible timing or phase noise: a minimum jitter in clocks, interferometers, resonators, or correlated measurements. The discriminator is a noise floor that remains after isolating all known noise sources and that scales with the relevant variables in a way predicted by a discreteness model.

These discriminators also apply to “becoming” theories that require a preferred present. A globally advancing “now” typically demands a preferred slicing, and preferred slicing tends to make itself visible as Lorentz-stressing structure unless carefully hidden.

14.3 Unitarity, reversibility, and resource-bounded physics

Not all discriminators are about symmetry breaking. Some are about whether evolution is fundamentally information-preserving, and whether the universe behaves like it has hard resource limits beyond ordinary locality.

Three clusters matter.

- Unitarity and information preservation

A broad class of physical descriptions treat closed-system evolution as reversible in principle: information is not destroyed, only redistributed. A rival class posits fundamental loss: information genuinely disappears, not merely becomes practically inaccessible. The discriminator is whether there exist controlled scenarios where information loss is irreducible even in principle, rather than merely hard to retrieve.

- Reversibility versus effective irreversibility

Thermodynamic irreversibility can emerge even when micro-evolution is reversible. If a proposal claims irreversibility is fundamental, it must predict a deviation from expected reversibility signatures in regimes engineered to minimize coarse-graining and environmental leakage. The discriminator is a clean separation between “we cannot reverse it because we lack control” and “no one can reverse it because the law forbids it.”

- Resource-bounded physics

If the universe is computational in a strong sense, there may be hard limits on

information throughput, state precision, or coordination depth that are stricter than ordinary finite signal speed. Such limits would appear as absolute ceilings on achievable correlations, synchronization quality, or error-correction performance, even in idealized isolation. The discriminator is not that experiments are difficult, but that there exists a principled bound that no amount of engineering can surpass.

These tests also help separate “iteration as metaphor” from “iteration as physics.” If iteration implies bounded update bandwidth per region per tick, the bound should have consequences for achievable correlations and control. If no bound manifests, iteration may be nothing over and above a reparameterization.

14.4 When theories are empirically equivalent but metaphysically distinct

Empirical equivalence is where many debates become permanently noisy. Two frameworks can disagree about what is fundamental while agreeing on every measurable prediction. In that case, physics does not adjudicate by data alone.

The paper treats this as a feature, not a failure. It implies a needed discipline in claims:

- If two descriptions are connected by a parameter change, a clock choice, or a gauge-like transformation that preserves all observable correlations, then they are the same theory in physics terms.
- If the only difference is narrative (time is “real flow” versus “coordinate”), and the narrative adds no coupling to observables, then the difference is interpretive, not experimental.
- If a metaphysical distinction is retained, it should be presented as a stance with virtues and costs (simplicity, explanatory unity, compatibility with other principles), not as a discovered fact.

This section’s practical rule is: never present a metaphysical add-on as if it were an empirical inference unless it produces an invariant, measurable discriminator.

14.5 The “minimal claim” ladder: from operational order to ontology

To keep discussion honest, the paper uses a ladder of claims. Moving up the ladder increases metaphysical commitment and decreases automatic empirical determinacy. Each rung is legitimate, but each requires different standards of justification.

- Level 0: Ordering as bookkeeping
“Time” or “iteration” is a parameter we use to label correlations. No claim that the label is itself physical.

- Level 1: Ordering from clocks
Time is defined operationally by a chosen clock subsystem. Physical content is in correlations relative to that clock.
- Level 2: Ordering from causality
The primitive ordering is influence (partial order). Global simultaneity is not assumed.
- Level 3: Arrow from thermodynamics and records
Directionality arises from low-entropy conditions and the accumulation of stable traces. This explains memory and prediction asymmetry without adding a moving present.
- Level 4: Discreteness with consequences
There is a true tick, event atom, or granular structure that cannot be removed by reparameterization and that implies at least one testable residue (symmetry stress, dispersion, anisotropy, noise floor, or principled bounds).
- Level 5: Becoming as physical structure
There exists an objective present or passage that adds structure beyond causal order and entropy/records, and that couples to observables (often via a preferred foliation or a fundamental asymmetry).

The ladder's purpose is not to forbid higher rungs. It is to prevent category slips: treating Level 0 or Level 1 language as if it implied Level 4 or Level 5 metaphysics, or treating Level 5 intuitions as if they were already Level 4 physics.

With these discriminators in place, the rest of the paper can compare "time," "iteration," and the other orderings without talking past itself: re-labelings are identified as re-labelings, clock choices as clock choices, and genuinely new physical claims are forced to state what would count as evidence for or against them.

15. Synthesis: A Map of Options

This paper's core claim is methodological: "time vs iteration" is not a well-posed fork until you specify what structure you are committing to and what would differ operationally if that structure were absent. We have now surveyed multiple candidates for "ordering": coordinate parameters, becoming/flow, discrete ticks, computational update rules, relational clocks, causal partial orders, entropy/records arrows, record-constraint ontologies, gauge/reparameterization freedom, phase cycles, scale flows, inferential/computational depth, and agency-centered update sequences. The task of synthesis is to turn that list into a map: a way to see which commitments imply which structures, which frameworks can coexist, and where "time" and "iteration" land inside the larger space.

15.1 A decision tree: which commitments imply which structures

A decision tree here is not a pop quiz. It is a dependency graph: if you affirm X, what else must you also accept (or explicitly reject) to remain coherent?

Start with the first branching question:

- Are you making a claim about *labels* or about *structure*?

If labels only, you are in the reparameterization/gauge zone. “Time” and “iteration” are merely coordinate choices until you add invariant structure that survives relabeling.

If structure, the next question is:

- Is your primitive ordering *global* or *local*?

If global, you are implicitly accepting something like a universal slicing (a global “now” or global tick). That pushes you toward a preferred frame/foliation risk and forces you to explain why global structure does not show up as symmetry-breaking residues.

If local, you naturally land in partial order territory: causality, local clocks, local record formation, and agent-local update sequences.

Then ask:

- Is your ordering *total* or *partial*?

Total order means every pair of events is comparable. Partial order means incomparability is allowed. Causal order is naturally partial. A global iteration counter tends toward total order unless carefully restricted.

Then ask:

- Is your ordering *reversible* or *arrowed*?

Reversible ordering is compatible with a time parameter without a direction. Arrowed ordering introduces asymmetry: entropy/records, fundamental irreversibility, complexity gradients, or learning arrows. If you want arrow without becoming, entropy/records is the standard route. If you want becoming, you owe additional structure beyond the arrow.

Then ask:

- Is your ordering *continuous* or *discrete*?

Discrete can be either mere discretization (relabeling, sampling) or fundamental (true tick). Fundamental discreteness implies non-removable scale and potential residues: dispersion, anisotropy, noise floors, or bounded information flow.

Finally ask:

- Is your ordering *ontic* or *epistemic*?

Ontic means it is part of the world's structure independent of observers (causal reachability, fundamental tick). Epistemic means it is tied to knowledge, inference, or agency (belief updates, derivability order). Epistemic orderings are valuable but must not be smuggled into claims about the world unless anchored to physical mechanisms (records, resource bounds, causal constraints).

In short, the decision tree says: once you decide (label vs structure), (global vs local), (total vs partial), (reversible vs arrowed), (continuous vs discrete), (ontic vs epistemic), your “time vs iteration” preference becomes a small downstream detail.

15.2 Compatibility chart: which frameworks can coexist

Most of the frameworks surveyed are not mutually exclusive. They can stack, provided you keep levels distinct.

Highly compatible combinations:

- Relational time + causal order
Causal structure supplies influence constraints; relational clocks supply operational parameterization inside that structure.
- Entropy arrow + records ontology
The arrow explains why records persist and accumulate; records explain why “past” is inferable from present traces. These are natural allies.
- Phase ordering + linear/arrow ordering
Cycles can be embedded within a larger directed structure: a fast phase variable plus slow drift plus record accumulation.
- Scale ordering + any event-ordering
Scale flow organizes descriptions; it does not compete with event-ordering. It complements time/causality by explaining universality and stability across levels.
- Inference depth + records/agency
Computational depth aligns naturally with record accumulation (“historical load”) and with learning dynamics in agents.

Potentially compatible, but tension-prone:

- Fundamental global tick + Lorentz-respecting causal structure
Possible if the tick is hidden by emergent symmetry, but it carries a heavy explanatory burden and invites symmetry-breaking discriminators.
- Becoming/flow + standard relativistic structure
Possible if “flow” is local or tied to a field/foliation that is empirically suppressed, but it must avoid either regress (meta-time) or epiphenomenalism (no observable coupling).

Often incompatible unless heavily engineered:

- Strong global present + no preferred structure
An objective global “now” typically implies a preferred foliation. If you forbid preferred structure, global present becomes difficult to make physical.
- Fundamental irreversibility + strict reversibility expectations
If you claim fundamental information loss, you must specify where it lives and why it does not contradict the success of reversible micro-descriptions in controlled regimes.

So the compatibility chart’s moral is pluralism with boundaries: many orderings can coexist as layered structures, but certain combinations force you to pay a symmetry or testability cost.

15.3 Where “time” and “iteration” sit inside the larger space

With the map in view, “time” and “iteration” become two special cases among many.

- Time as coordinate
Sits in the label/gauge zone unless anchored to clocks. It is indispensable as a tool, but not automatically ontic.
- Time as becoming
Sits at the high-commitment end: it adds structure beyond correlations, beyond causal order, beyond entropy arrows. To be physics, it must couple to observables.
- Iteration as reparameterization
Sits in the same label/gauge zone as coordinate time. It is a different index, not a different world.
- Iteration as clock counting
Sits inside relational time: n is the discrete readout of a physical periodic process plus record-keeping.

- Iteration as true tick
Sits in the discreteness-with-consequences region: it claims a non-removable grain, potentially with symmetry or noise-floor residues.
- Iteration as computation with bounds
Sits in the resource-bounded physics region: it claims hard limits on information flow, coordination depth, or precision beyond ordinary locality.

So the synthesis is blunt: “time vs iteration” is only a real fork in the tree when “iteration” means true tick or resource-bounded update, or when “time” means becoming with physical coupling. Otherwise, both are usually just different languages for the same correlation structure.

15.4 A pragmatic stance: disciplined pluralism and falsifiable edges

The stance recommended by this paper is pragmatic but not lax: disciplined pluralism with falsifiable edges.

Disciplined pluralism means:

- Use whatever ordering is appropriate to the task: clock time for laboratory prediction, causal order for influence constraints, entropy/records for directionality and memory, phase for synchronization, scale for universality, inference/agency for cognition and learning.
- Treat “time” and “iteration” as representational choices unless you explicitly add structure.
- Keep metaphysical preferences explicit as preferences, not as empirical conclusions.

Falsifiable edges means:

- If you claim a true tick, specify what residue it leaves: symmetry stress, dispersion, anisotropy, irreducible noise floors, or bounded correlation/throughput limits.
- If you claim becoming as physical, specify the added structure and how it couples: preferred foliation, new field variable, time-asymmetric micro-dynamics, or an operational signature that cannot be reparameterized away.
- If you claim computational resource bounds, specify the bound and what experimental domain would be sensitive to it.

This stance respects both the power and the limits of physics. Physics can tell us what structures are required to reproduce observed invariants. It can constrain additional structure by symmetry tests and precision measurements. But when multiple stories are empirically

equivalent, the honest move is not to declare a winner by rhetoric. It is to say: here is the shared operational core, here are the alternative ontological readings, and here are the edges where they would diverge if they ever become more than interpretation.

That, ultimately, is the map: a way to keep “time vs iteration” from being too narrow, too loaded, or too easy—and to keep our metaphysics tethered to what could, even in principle, matter to the world.

16. Conclusion: How to Speak About Order Without Overclaiming

This paper began with a narrow-seeming question—are we in “time” or “iteration”?—and treated it as a diagnostic. The diagnosis is that the question is often underspecified: many disputes are really about labels, clock choices, or narrative intuitions. Once we enlarge the space of possibilities, we find a spectrum of orderings that physics can use coherently: coordinate parameters, relational clocks, causal partial orders, entropy arrows, record constraints, phase cycles, scale flows, inferential depth, and agent-centered update sequences. Each orders something real, but not all order the same kind of thing. The discipline, then, is not to pick a single word and crown it fundamental, but to speak precisely about *which ordering relation is being claimed, what structure it adds, and what would count as evidence that it is more than a re-description*.

16.1 Restating the central result: many orderings, few discriminators

The central result is simple to state and surprisingly hard to internalize:

- There are many internally consistent ways to define “before/after” or “order.”
- Only a small subset of differences between these ways produce distinct empirical predictions.

Most of the time/iteration debate collapses into one of two low-commitment zones:

- **Reparameterization/gauge:** “time” or “iteration” is a label that can be changed without changing observables.
- **Clock choice/relational time:** “time” is a variable read off from a physical subsystem, and different clocks give different but compatible coordinatizations.

Where genuine discriminators appear is at the “falsifiable edges”:

- A **true tick** that implies a non-removable granularity with detectable residues (symmetry stress, dispersion, anisotropy, noise floors, throughput bounds).

- A **physical becoming** that adds structure beyond causality and entropy/records and couples to observables (preferred foliation, time-asymmetric micro-dynamics, an operational “presentness” variable that does work).
- A **resource-bounded computation** claim that produces principled, testable limits beyond ordinary locality.

Everything else—entropy arrows, records, phase cycles, scale flow, inferential depth, agency updates—can be real and powerful while remaining compatible with multiple underlying ontologies. They often explain *why experience looks ordered* without forcing a single primitive ordering variable.

So the conclusion is not “nothing is real.” It is: real structure lives in invariants and constraints; the words we use to narrate that structure must not outrun the discriminators.

16.2 Practical takeaway for theorists and communicators

For theorists, the practical takeaway is a design rule:

- Whenever you propose a new ordering primitive, immediately state what is invariant and what changes.
- Specify whether your ordering is global or local, total or partial, continuous or discrete, reversible or arrowed, ontic or epistemic.
- Identify the exact point where your proposal becomes more than a reparameterization or clock choice.

A good proposal should answer four questions in one paragraph:

1. What is the primitive relation (correlation, causality, tick, record constraint, phase, scale, derivation, agency)?
2. What additional structure does it posit beyond standard descriptions?
3. What would differ in observable correlation structure if that additional structure were absent?
4. What existing constraints or potential experiments bear on those differences?

For communicators—writers, teachers, public explainers—the takeaway is a hygiene rule:

- Stop using “time is an illusion” and “the universe updates like a computer” as if they were settled physical conclusions.

- Replace them with claims at the appropriate rung of the ladder: “time is operationally defined by clocks,” “many parameter choices are gauge-like,” “causal order is more robust than simultaneity,” “entropy and records produce an arrow,” “discreteness would require non-removable residues,” “becoming would require extra structure.”

When discussing experience (“flow feels unavoidable”), separate three layers:

- **Operational layer:** what clocks, signals, and records measure.
- **Structural layer:** what invariants and constraints physics commits to.
- **Interface layer:** what bounded agents experience (memory, updating, control loops).

This separation reduces confusion and prevents “experience language” from being misread as “new physics,” while still treating experience as a serious datum for cognitive science and for any future theory that tries to connect mind and world.

16.3 The open frontier: new tests, new formalisms, cleaner language

The open frontier is not merely “more data.” It is cleaner conceptual engineering.

Three frontiers stand out.

- **New tests at the falsifiable edges**
The most valuable experiments will be those that can distinguish fundamental discreteness or preferred structure from mere noise and modeling artifacts: tighter symmetry stress tests, long-baseline propagation constraints, cross-platform noise-floor comparisons, and sharper bounds on any hypothesized resource ceilings. Even null results matter because they prune entire families of “global tick” and “preferred present” stories.
- **New formalisms that factor ordering cleanly**
We need frameworks where different orderings can be expressed as different layers without confusion: causal structure separated from metric structure; clock choice separated from invariants; entropy/records arrows separated from parameters; agency update sequences separated from ontic causality. The aim is not one grand metaphysical declaration, but a compositional language in which multiple orderings can coexist without being conflated.
- **Cleaner language for public and cross-disciplinary work**
The biggest conceptual progress often comes from refusing to let one overloaded word do many jobs. “Time” is such a word. “Iteration” is becoming such a word. The frontier is a vocabulary that names the distinct structures explicitly: parameter, clock variable, causal precedence, record constraint, entropy gradient, phase, scale flow, derivation

depth, decision update. When the vocabulary is clean, overclaiming becomes harder, and real novelty becomes easier to spot.

The paper's closing posture is therefore both modest and ambitious. Modest: physics rarely forces a single metaphysical reading when multiple parameterizations are empirically equivalent. Ambitious: the space of orderings is richer than "time vs iteration," and the right kind of rigor—about invariants, discriminators, and layered structure—can make that richness a source of progress rather than confusion.

Order is not one thing. It is a family of structures. Speak about it with the family resemblance in view, and with the falsifiable edges clearly marked. That is how we stay honest, stay curious, and leave room for the next genuine step: a theory that does not merely rename time, but explains why so many different orderings work so well—and where they finally diverge.

17. References

Note: The entries below are grouped by theme to match the paper's taxonomy. They include foundational sources, widely used reviews, and representative technical treatments.

17.1 Relational and emergent time

Page, D. N., and Wootters, W. K. (1983). "Evolution without evolution: Dynamics described by stationary observables." *Physical Review D*, 27(12), 2885–2892.

Rovelli, C. (1990). "Quantum mechanics without time: A model." *Physical Review D*, 42(8), 2638–2646.

Rovelli, C. (1991). "Time in quantum gravity: An hypothesis." *Physical Review D*, 43(2), 442–456.

Connes, A., and Rovelli, C. (1994). "Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories." *Classical and Quantum Gravity*, 11(12), 2899–2918.

Barbour, J. (1994). "The timelessness of quantum gravity: I. The evidence from the classical theory." *Classical and Quantum Gravity*, 11(12), 2853–2873.

Barbour, J. (1999). *The End of Time*. Oxford University Press.

Gambini, R., Porto, R. A., and Pullin, J. (2004). "Fundamental decoherence from quantum gravity: A pedagogical review." *General Relativity and Gravitation*, 39, 1143–1156.

Gambini, R., and Pullin, J. (2009). "Relational physics with real rods and clocks and the measurement problem of quantum mechanics." *Foundations of Physics*, 39, 1143–1156.

Oreshkov, O., Costa, F., and Brukner, C. (2012). "Quantum correlations with no causal order." *Nature Communications*, 3, 1092.

17.2 Causal set and causal order approaches

Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. (1987). "Space-time as a causal set." *Physical Review Letters*, 59(5), 521–524.

Rideout, D. P., and Sorkin, R. D. (2000). "A classical sequential growth dynamics for causal sets." *Physical Review D*, 61(2), 024002.

Sorkin, R. D. (2003). "Causal sets: Discrete gravity." In *Lectures on Quantum Gravity* (Series of the Centro De Estudios Científicos), Springer.

Henson, J. (2006). "The causal set approach to quantum gravity." In *Approaches to Quantum Gravity*, Cambridge University Press.

Dowker, F. (2005). "Causal sets and the deep structure of spacetime." In *100 Years of Relativity: Space-Time Structure*, World Scientific.

Benincasa, D. M. T., and Dowker, F. (2010). "The scalar curvature of a causal set." *Physical Review Letters*, 104(18), 181301.

Surya, S. (2019). "The causal set approach to quantum gravity." *Living Reviews in Relativity*, 22,

5.

Malament, D. B. (1977). "The class of continuous timelike curves determines the topology of spacetime." *Journal of Mathematical Physics*, 18(7), 1399–1404.

17.3 Thermodynamic arrow and low-entropy boundary work

Boltzmann, L. (1896–1898). *Lectures on Gas Theory*. (English translations, Dover and other editions).

Loschmidt, J. (1876). "On the state of thermal equilibrium of a system of bodies with regard to the force of gravity." (Foundational reversibility objection; historical source).

Zermelo, E. (1896). "On the mechanical explanation of irreversible processes." (Historical recurrence objection; sources in collected editions).

Reichenbach, H. (1956). *The Direction of Time*. University of California Press.

Penrose, R. (1979). "Singularities and time-asymmetry." In *General Relativity: An Einstein Centenary Survey*, Cambridge University Press.

Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press.

Lebowitz, J. L. (1993). "Boltzmann's entropy and time's arrow." *Physics Today*, 46(9), 32–38.

Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford University Press.

Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.

Carroll, S. M. (2010). *From Eternity to Here: The Quest for the Ultimate Theory of Time*. Dutton.

17.4 Records, decoherence, and information-theoretic treatments

Zurek, W. H. (2003). "Decoherence, einselection, and the quantum origins of the classical." *Reviews of Modern Physics*, 75(3), 715–775.

Zurek, W. H. (2009). "Quantum Darwinism." *Nature Physics*, 5, 181–188.

Joos, E., Zeh, H. D., Kiefer, C., Giulini, D., Kupsch, J., and Stamatescu, I.-O. (2003). *Decoherence and the Appearance of a Classical World in Quantum Theory*. Springer.

Schlosshauer, M. (2007). *Decoherence and the Quantum-to-Classical Transition*. Springer.

Gell-Mann, M., and Hartle, J. B. (1990). "Quantum mechanics in the light of quantum cosmology." In *Complexity, Entropy, and the Physics of Information*, Addison-Wesley.

Griffiths, R. B. (1984). "Consistent histories and the interpretation of quantum mechanics." *Journal of Statistical Physics*, 36, 219–272.

Omnès, R. (1992). "Consistent interpretations of quantum mechanics." *Reviews of Modern Physics*, 64(2), 339–382.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press.

Shannon, C. E. (1948). "A mathematical theory of communication." *Bell System Technical Journal*, 27, 379–423 and 623–656.

Landauer, R. (1961). "Irreversibility and heat generation in the computing process." *IBM Journal*

of Research and Development, 5(3), 183–191.

Bennett, C. H. (1982). “The thermodynamics of computation: A review.” *International Journal of Theoretical Physics*, 21(12), 905–940.

Jaynes, E. T. (1957). “Information theory and statistical mechanics.” *Physical Review*, 106, 620–630 and 108, 171–190.

17.5 Gauge time, parameterization invariance, and canonical gravity

Dirac, P. A. M. (1958). “The theory of gravitation in Hamiltonian form.” *Proceedings of the Royal Society A*, 246, 333–343.

Arnowitt, R., Deser, S., and Misner, C. W. (1962). “The dynamics of general relativity.” In *Gravitation: An Introduction to Current Research*, Wiley.

Wheeler, J. A. (1968). “Superspace and the nature of quantum geometrodynamics.” In *Battelle Rencontres*, Benjamin.

DeWitt, B. S. (1967). “Quantum theory of gravity. I. The canonical theory.” *Physical Review*, 160, 1113–1148.

Kuchař, K. V. (1992). “Time and interpretations of quantum gravity.” In *Proceedings of the 4th Canadian Conference on General Relativity and Relativistic Astrophysics*, World Scientific.

Isham, C. J. (1993). “Canonical quantum gravity and the problem of time.” In *Integrable Systems, Quantum Groups, and Quantum Field Theories*, Springer.

Henneaux, M., and Teitelboim, C. (1992). *Quantization of Gauge Systems*. Princeton University Press.

Rovelli, C. (2004). *Quantum Gravity*. Cambridge University Press.

Thiemann, T. (2007). *Modern Canonical Quantum General Relativity*. Cambridge University Press.

Barbour, J., Foster, B. Z., and Ó Murchadha, N. (2002). “Relativity without relativity.” *Classical and Quantum Gravity*, 19(12), 3217–3248.

17.6 Discreteness tests, Lorentz invariance constraints, dispersion searches

Mattingly, D. (2005). “Modern tests of Lorentz invariance.” *Living Reviews in Relativity*, 8, 5.

Kostelecký, V. A. (2004). “Gravity, Lorentz violation, and the standard model.” *Physical Review D*, 69, 105009.

Kostelecký, V. A., and Russell, N. (2011–present). “Data tables for Lorentz and CPT violation.” *Reviews of Modern Physics* (periodic updates and arXiv versions).

Coleman, S., and Glashow, S. L. (1999). “High-energy tests of Lorentz invariance.” *Physical Review D*, 59, 116008.

Jacobson, T., Liberati, S., and Mattingly, D. (2006). “Lorentz violation at high energy: Concepts, phenomena and astrophysical constraints.” *Annals of Physics*, 321(1), 150–196.

Liberati, S. (2013). “Tests of Lorentz invariance: A 2013 update.” *Classical and Quantum Gravity*,

30, 133001.

Amelino-Camelia, G. (2002). "Doubly-special relativity: First results and key open problems." *International Journal of Modern Physics D*, 11(10), 1643–1669.

Abdo, A. A., et al. (Fermi LAT and GBM Collaborations) (2009). "A limit on the variation of the speed of light arising from quantum gravity effects." *Nature*, 462, 331–334.

Ellis, J., Mavromatos, N. E., and Nanopoulos, D. V. (2008). "Quantum-gravity analysis of gamma-ray bursts using wavelets." *Physics Letters B*, 665, 412–417.

Amelino-Camelia, G., et al. (1998). "Tests of quantum gravity from observations of gamma-ray bursts." *Nature*, 393, 763–765.

17.7 Renormalization group and scale as ordering

Wilson, K. G. (1971). "Renormalization group and critical phenomena. I. Renormalization group and the Kadanoff scaling picture." *Physical Review B*, 4, 3174–3183.

Wilson, K. G., and Kogut, J. (1974). "The renormalization group and the epsilon expansion." *Physics Reports*, 12(2), 75–199.

Fisher, M. E. (1974). "The renormalization group in the theory of critical behavior." *Reviews of Modern Physics*, 46(4), 597–616.

Polchinski, J. (1984). "Renormalization and effective Lagrangians." *Nuclear Physics B*, 231(2), 269–295.

Weinberg, S. (1995). *The Quantum Theory of Fields, Vol. I*. Cambridge University Press.

Weinberg, S. (1996). *The Quantum Theory of Fields, Vol. II*. Cambridge University Press.

Goldenfeld, N. (1992). *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley.

Cardy, J. (1996). *Scaling and Renormalization in Statistical Physics*. Cambridge University Press.

Zinn-Justin, J. (2002). *Quantum Field Theory and Critical Phenomena*. Oxford University Press.

17.8 Complexity, computational depth, and inference-based arrows

Kolmogorov, A. N. (1965). "Three approaches to the quantitative definition of information." *Problems of Information Transmission*, 1, 1–7.

Solomonoff, R. J. (1964). "A formal theory of inductive inference." *Information and Control*, 7, 1–22 and 224–254.

Chaitin, G. J. (1975). "A theory of program size formally identical to information theory." *Journal of the ACM*, 22(3), 329–340.

Bennett, C. H. (1988). "Logical depth and physical complexity." In *The Universal Turing Machine: A Half-Century Survey*, Oxford University Press.

Crutchfield, J. P., and Young, K. (1989). "Inferring statistical complexity." *Physical Review Letters*, 63, 105–108.

Lloyd, S. (2000). "Ultimate physical limits to computation." *Nature*, 406, 1047–1054.

Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.

Aaronson, S. (2013). *Quantum Computing Since Democritus*. Cambridge University Press.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.

Friston, K. (2010). "The free-energy principle: A unified brain theory?" *Nature Reviews Neuroscience*, 11, 127–138.

17.9 Philosophy of time, flow, and operationalism

McTaggart, J. M. E. (1908). "The unreality of time." *Mind*, 17(68), 457–474.

Putnam, H. (1967). "Time and physical geometry." *Journal of Philosophy*, 64(8), 240–247.

Grünbaum, A. (1973). *Philosophical Problems of Space and Time* (2nd ed.). Reidel.

Sklar, L. (1974). *Space, Time, and Spacetime*. University of California Press.

Maudlin, T. (2002). "Remarks on the passing of time." *Proceedings of the Aristotelian Society*, 102, 237–252.

Mellor, D. H. (1998). *Real Time II*. Routledge.

Callender, C. (2017). *What Makes Time Special?* Oxford University Press.

Bridgman, P. W. (1927). *The Logic of Modern Physics*. Macmillan.

Einstein, A. (1905). "On the electrodynamics of moving bodies." *Annalen der Physik*, 17, 891–921.

Einstein, A. (1916). "The foundation of the general theory of relativity." *Annalen der Physik*, 49, 769–822.

